

inspur

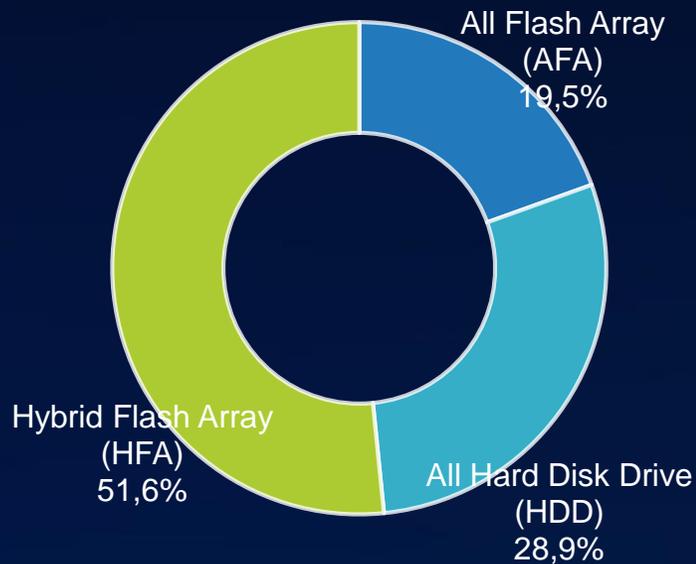
Cloud storage, intelligent use, advance new data

Why Inspur AFA?

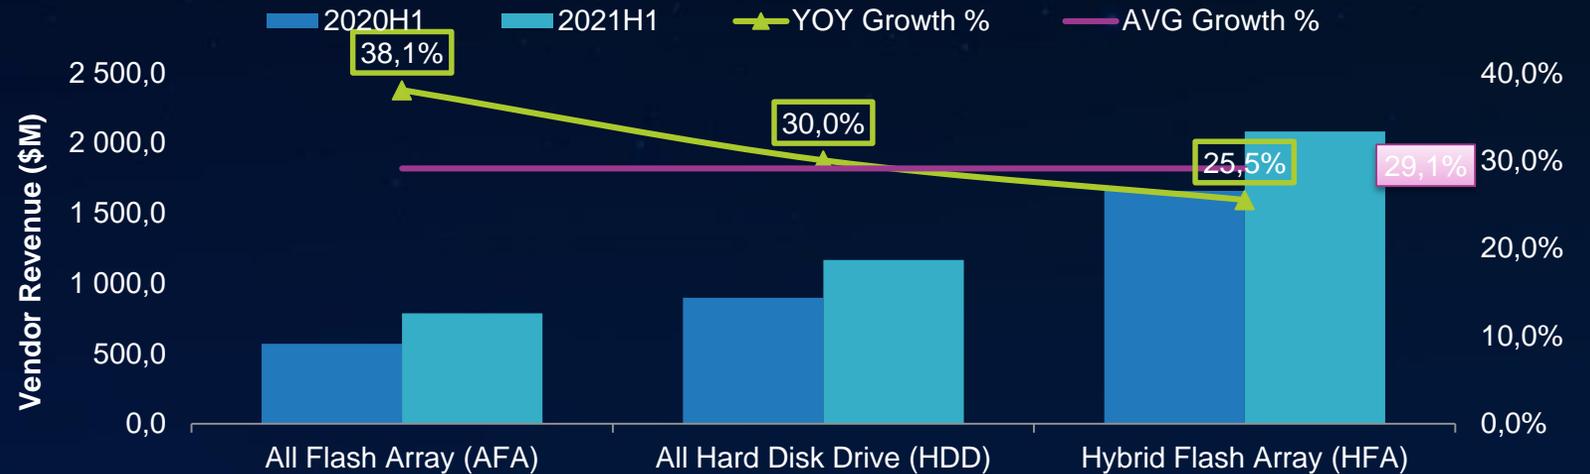
With the rapid development of full flash storage Become a trend to replace Hybrid Flash & pure HDD storage



**PRC Storage Market Overview
by Storage Array Type, 2021YTD**



PRC Flash Storage Market Overview, 2021YTD/2020YTD

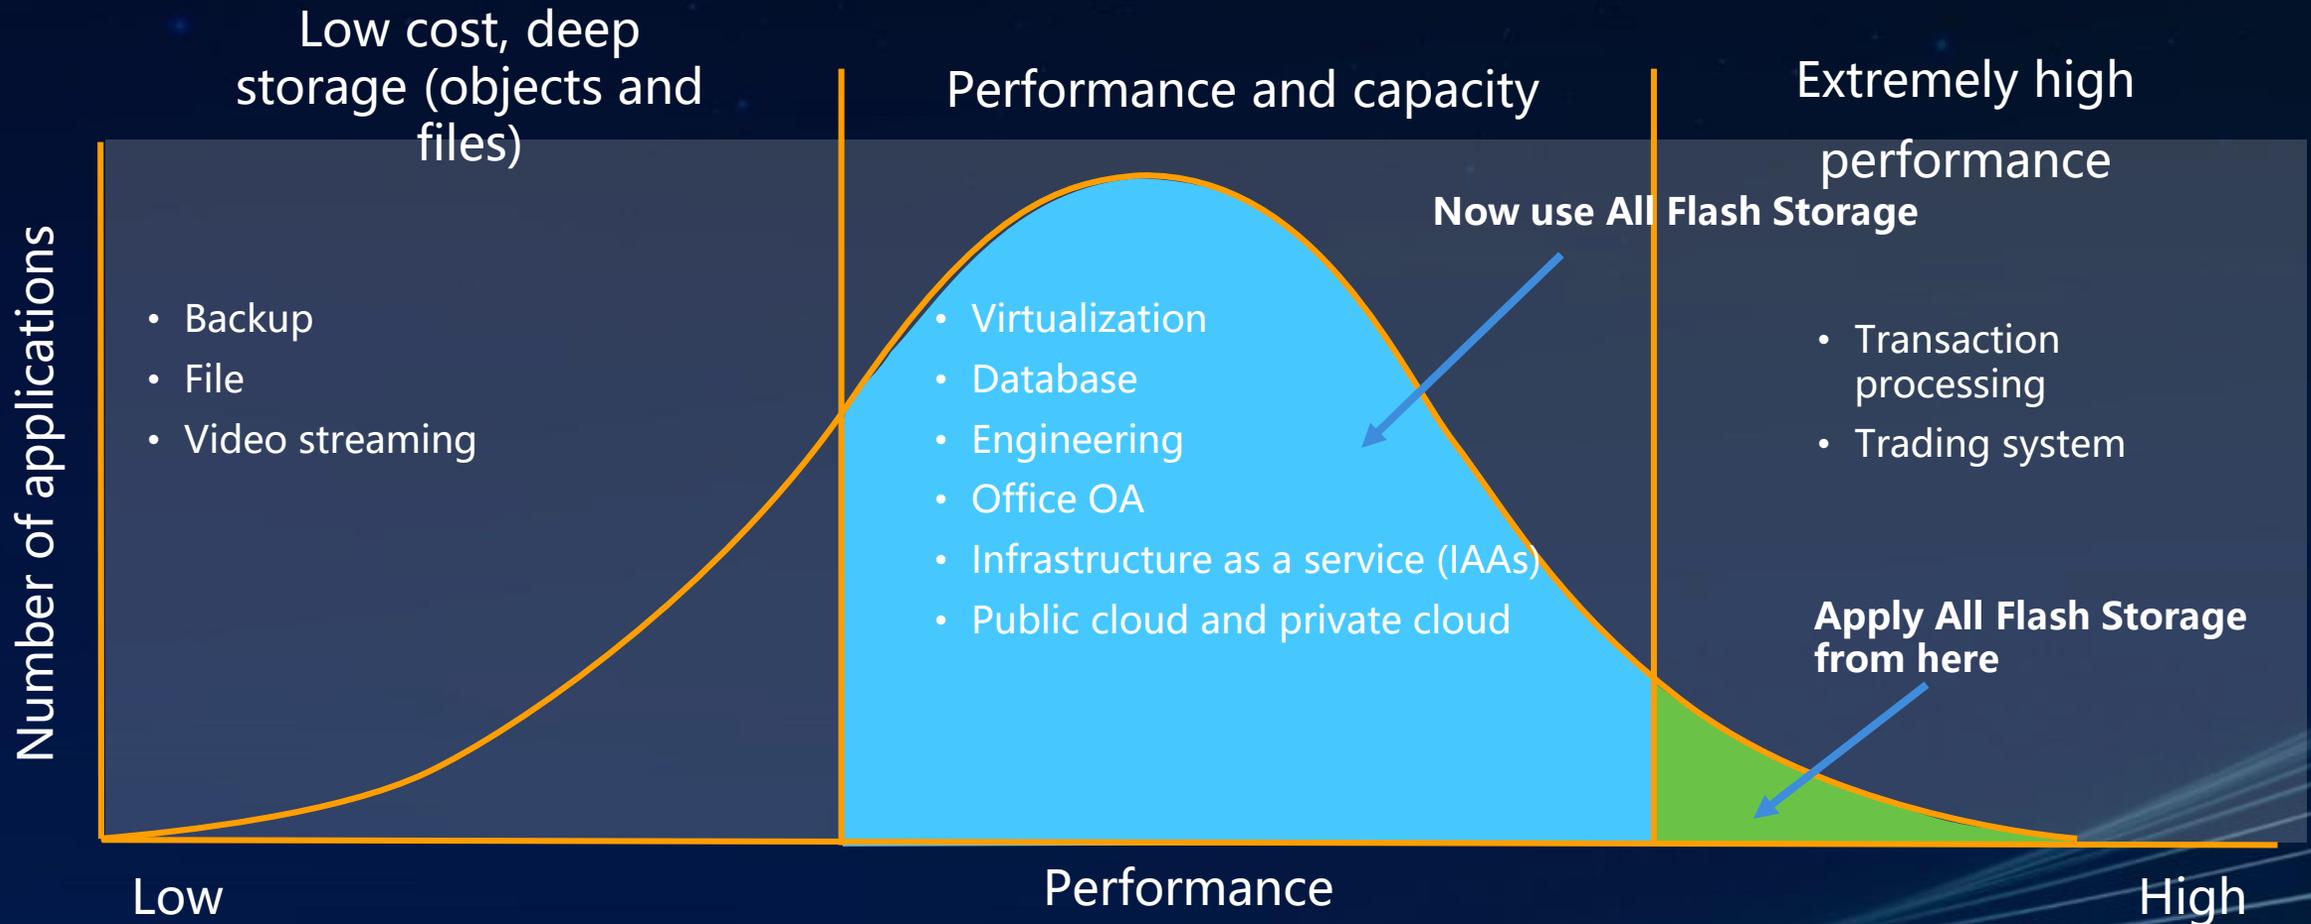


- China's all flash storage market is developing rapidly, with an annual growth of nearly 40% and a market share of about 19.5%
- Inspur ranks second in 2021Q2 / Q3 market share in the Chinese market
- In the global market, the overall proportion of full flash is about 40%, and the proportion of us and EMEA markets is higher

All flash storage transformation is at the right time



- The cost of flash memory is gradually approaching SAS disk, and can provide greater storage capacity and space density



All Flash Storage helps customers get better TCO



All Flash ≠ Expensive

HF5000G5 VS. Hybrid Flash products from other competitors:

Under IOPS fixed conditions:

Same grade Hybrid Flash Storage of competitive products: configured with 1.2T SAS 10K disk and RAID5

Configure 450 disks, 8K, 7:3 read-write ratio, random IO

The performance is 100000 IOPs

18 cabinets are required

HF5000G5: 20 1.92TB SSD disks, 8K, 7:3 read-write ratio, random IO

The performance is 300000 IOPs

TCO is only one tenth of the original storage:

1

10

| | | |
|---|----------|--|
| Computer room construction cost | 5 | 40000-60,000 RMB/m ² (including decoration, power supply, refrigeration, cabinets and supporting facilities, excluding civil and IT equipment) |
| Air conditioning power consumption | 0.000293 | Calculated according to the heat dissipation of the equipment, 1BTU/hour = 0.000293 KW |
| Commercial electricity price in Beijing | 0.781 | The price is ¥0.781/kw, reference website: http://www.bj.sgcc.com.cn/Contents/Channel_52/2009/1123/58255/content_58255.htm |
| Number of TCO years | 6 | |

| | 项目 | 原有 | 替换方案 |
|-------|---------------|----------|----------|
| 设备配置 | 型号 | AS5500G5 | HF5000G5 |
| | 可用容量(TB) | 450 | 30 |
| | 磁盘类型(GB) | 1229 | 1920 |
| | RAID级别 | RAID5 | RAID5 |
| | RAID5条带 | 9 | 9 |
| | 热备盘数 | 13 | 1 |
| 性能 | 磁盘总数量 | 436 | 20 |
| | IOPS(K) | 100 | 300 |
| 占地成本 | Latency(ms) | 5 | 1 |
| | 重量(KG) | 504 | 28 |
| | 占地面积(平方米) | 0.54 | 0.27 |
| 能耗成本 | 机房建设成本(万元) | 2.70 | 1.35 |
| | 功率(KW) | 13.5 | 0.75 |
| | 散热 (BTU/Hour) | 39957 | 2280 |
| | 6年设备耗电(千瓦时) | 709560 | 39420 |
| | 6年制冷耗电(千瓦时) | 615341 | 35112 |
| 维保 | 6年电费(万元) | 103.47 | 5.82 |
| | 每年维保(万元) | 1 | 1 |
| TCO合计 | 6年总计 (万元) | 112.17 | 13.17 |
| TCO节省 | 相对原方案节省 | 88.26% | N/A |

The normal maintenance cost is based on the 7-8% of the sales price at that time.

High-end All Flash Storage SPC-1™ Performance Chart No. 1 in the World



In August 2021, Inspur HF18000G5 storage reached 23 million IOPS, ranking the first in the world.

HF18000G5 lowest system latency
Performance No. 1 in the world

SPC-1IOPSTM 23,001,502
SPC-1Price-Performance \$375/KIOPSTM
SPC-1IOPSResponseTime 0.294ms
SPC-1OverallResponseTime 0.246ms



Data source: SPC-1 official website

Inspur Storage Product SPC-1™ Performance Chart

No. 1 in the World



In March, June and September, 2020, Inspur's 16-controller and 8-Controller unified storage products ranked No.1 in the world in single-control performance

Inspur mid-end and high-end storage products

Performance No. 1 in the world

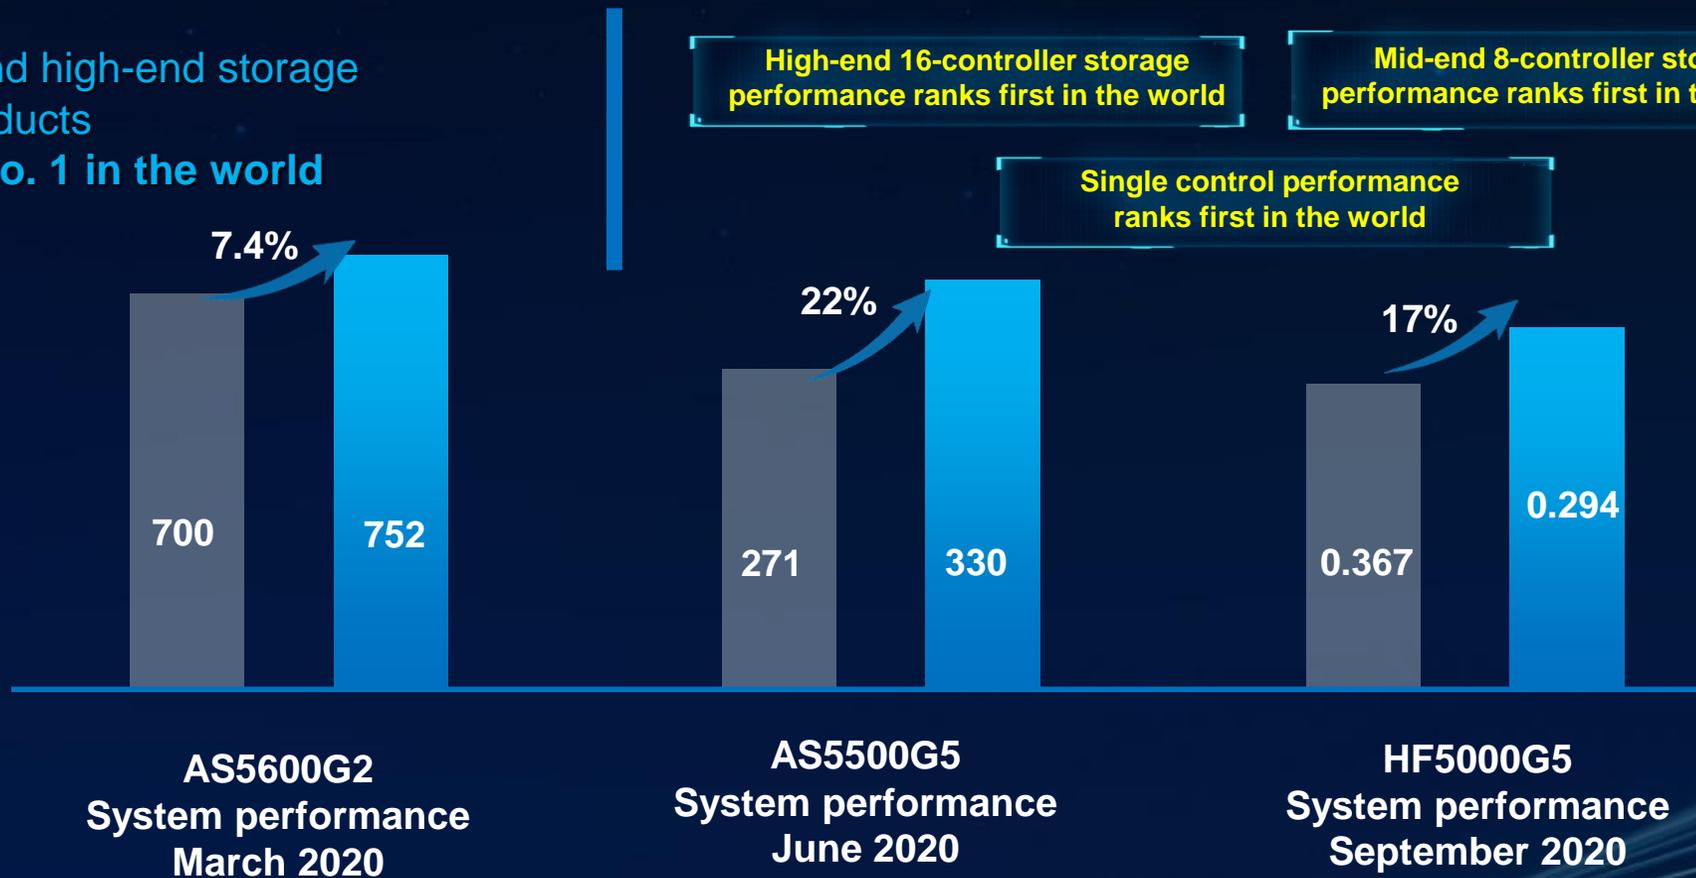
High-end 16-controller storage performance ranks first in the world

Mid-end 8-controller storage performance ranks first in the world

Single control performance ranks first in the world

单位: 万IOPS

Inspur
Former No.1



Data source: SPC-1 official website

Inspur all flash storage product family introduction

inspur

One set of architecture

Multiple media fusion

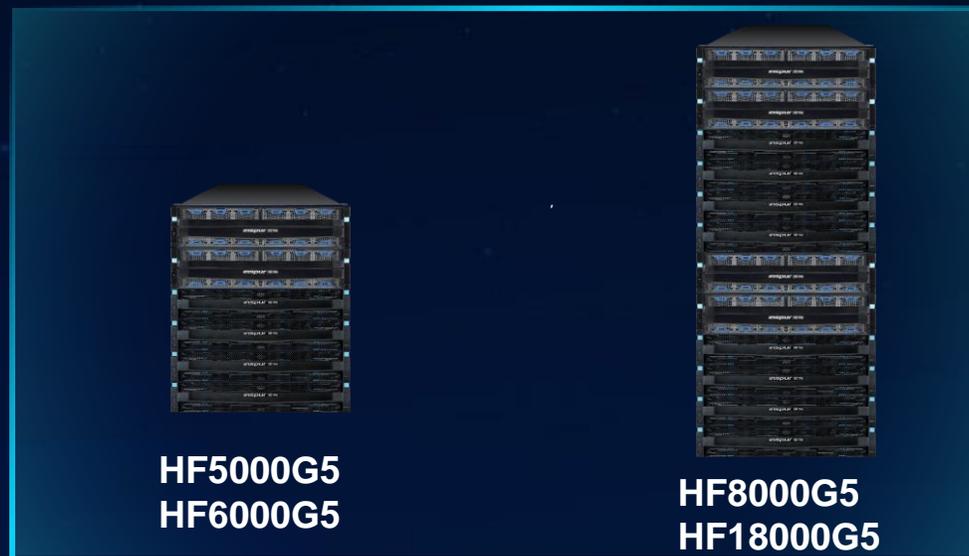
Maximum 32 controllers

Performance lossless
compression

AA, 3DC

WAN acceleration, **Cloud gateway**

Leading architecture



Fully redundant
components

Health management

Fault prediction

Automatic fault
repair

Multi control cache
copy

Fast data
reconstruction

**Comprehensive safety
and reliability design**

iTurbo 2.0

Multiple performance
accelerator engine

Performance **No.1** in the world

23 million IOPS

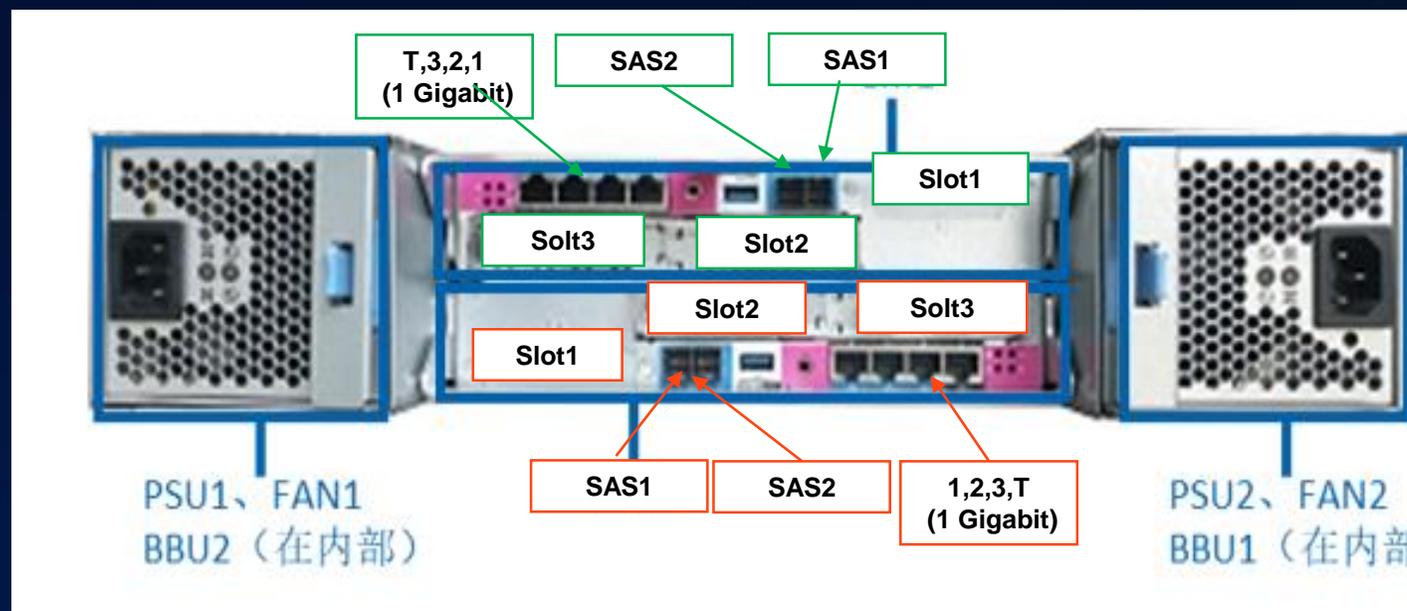
5 times less delay

0.2ms

Leading performance

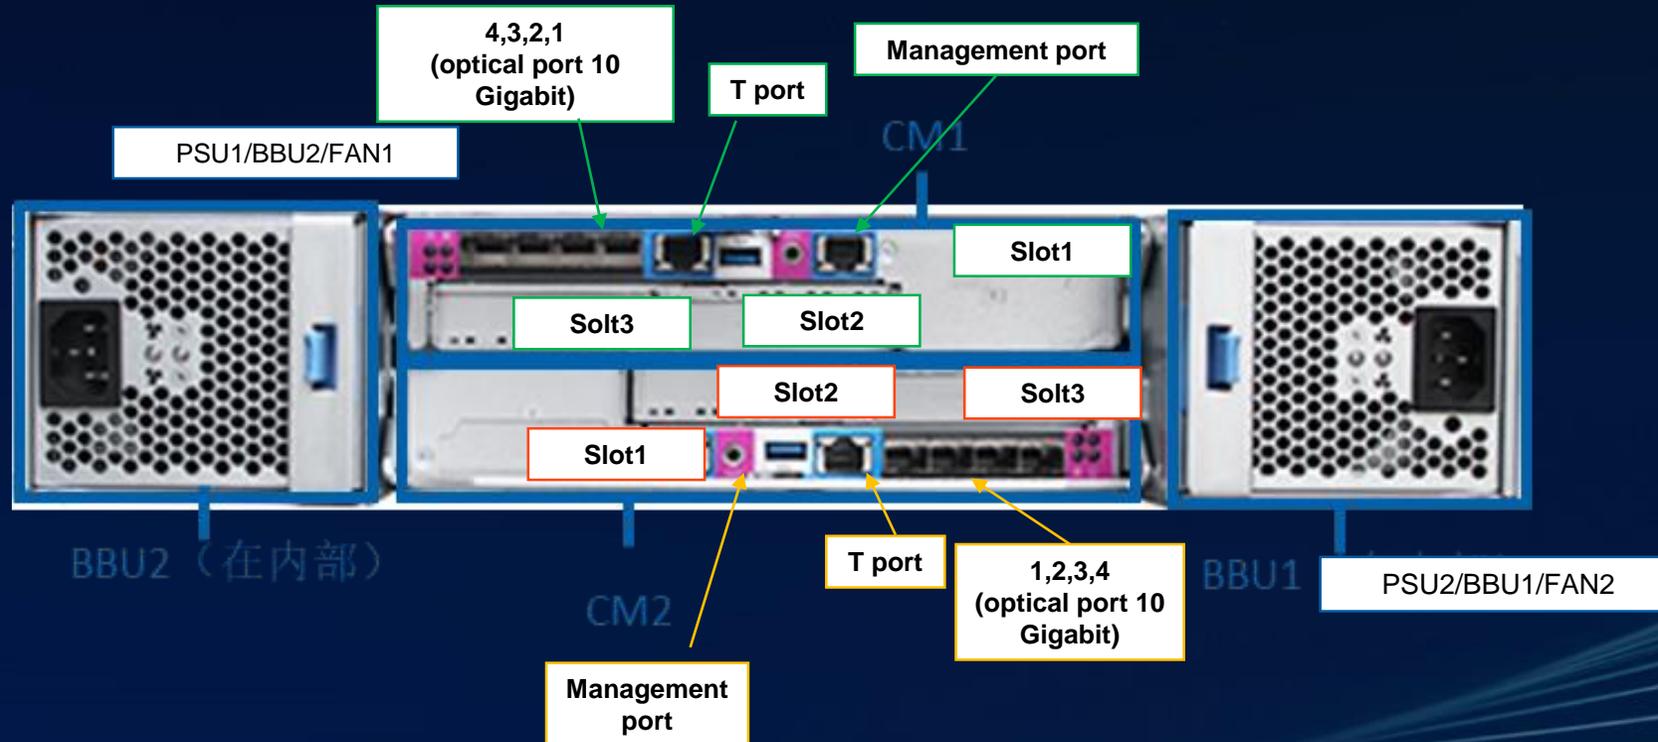
HF5000G5-MS25/HS25 Physical form

inspur



HF5000G5-MN25 Physical form

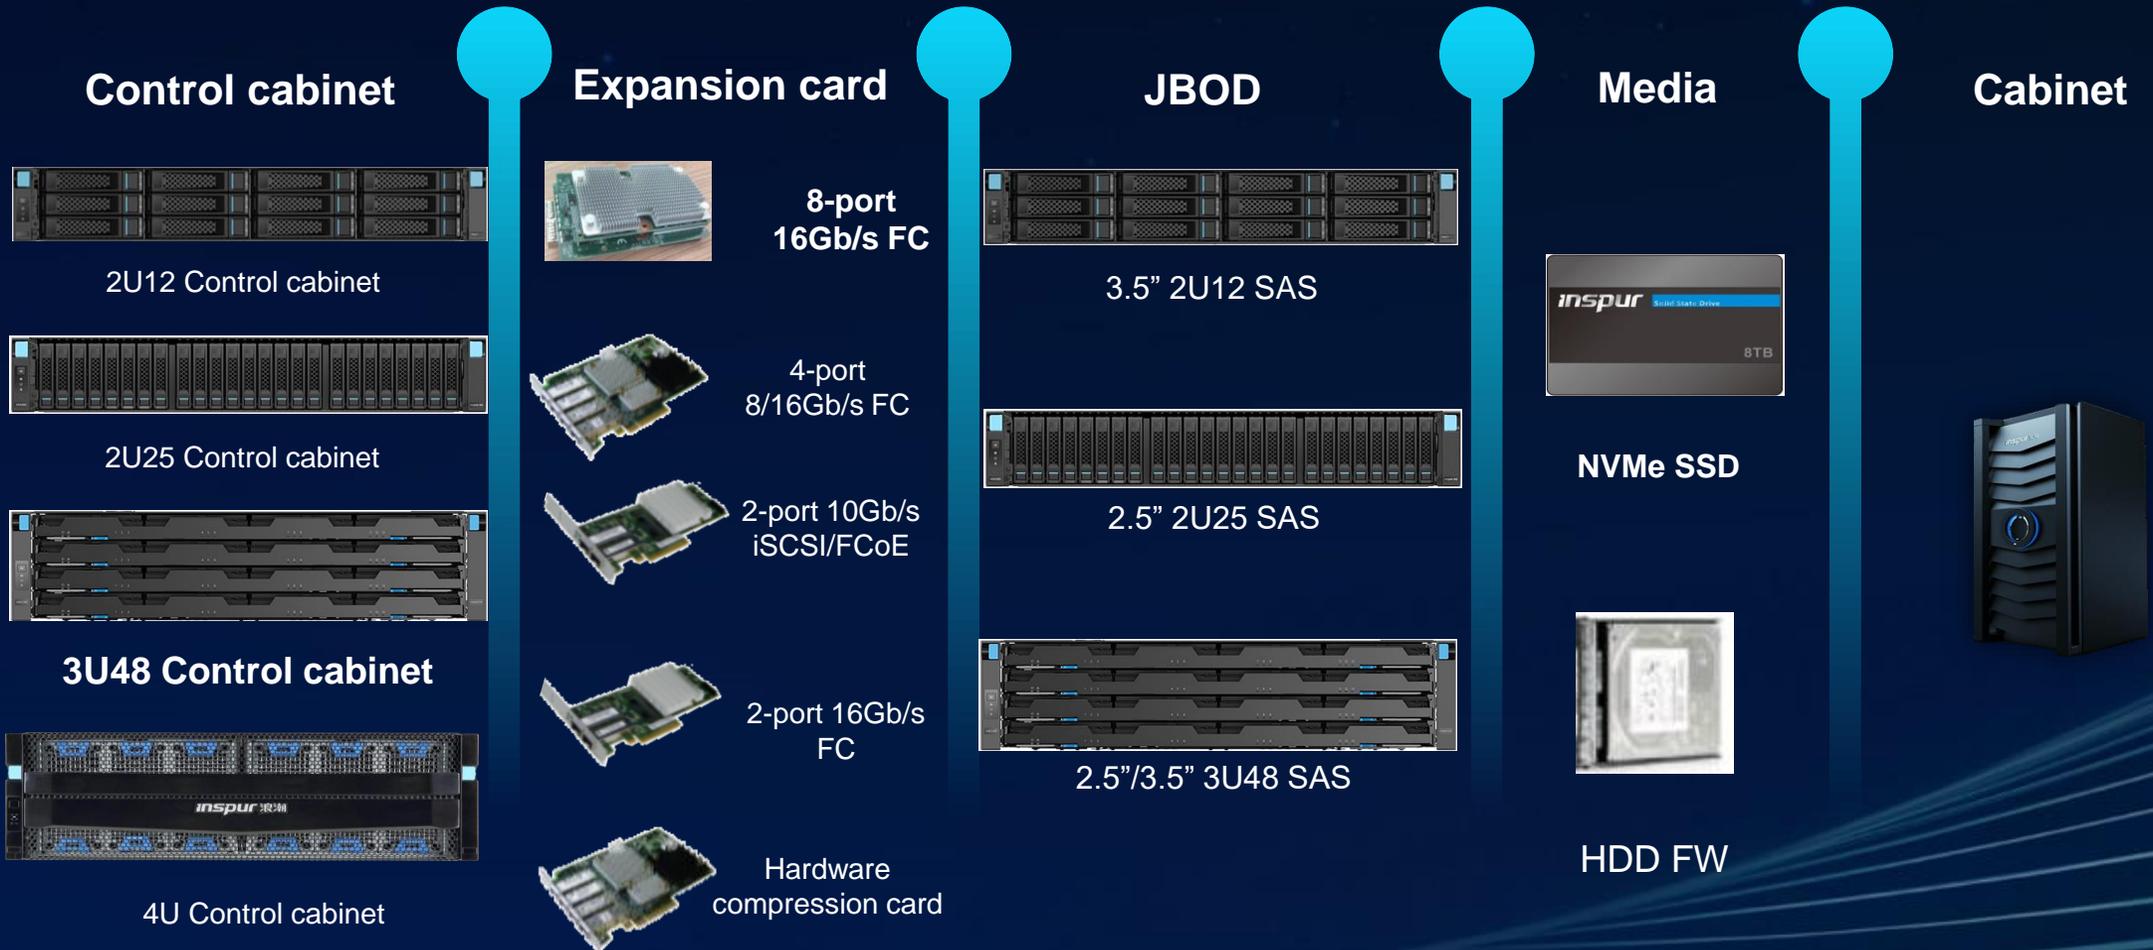
inspur



Autonomous and controllable unified storage hardware platform

inspur

- Self research and design of storage system software and hardware: host, expansion card, JBOD, SSD
- All flash storage system with the largest cache, 8-port FC card with the largest number of ports, self-developed NVMe, SSD FW

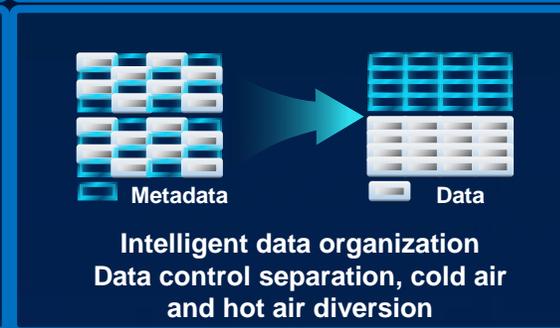
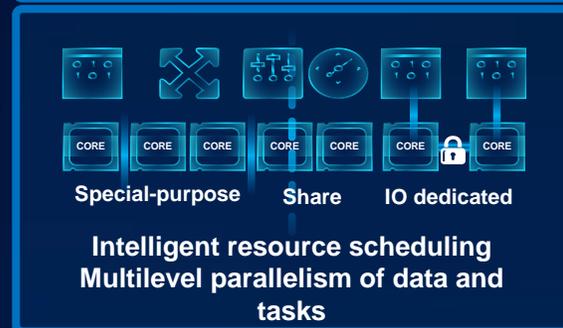
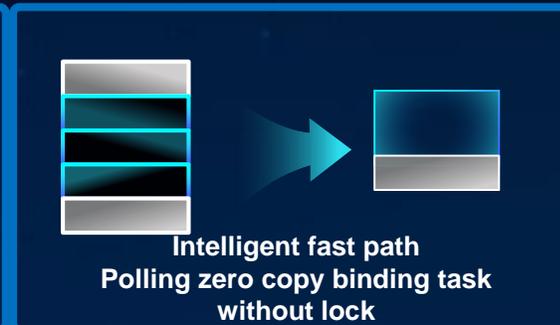
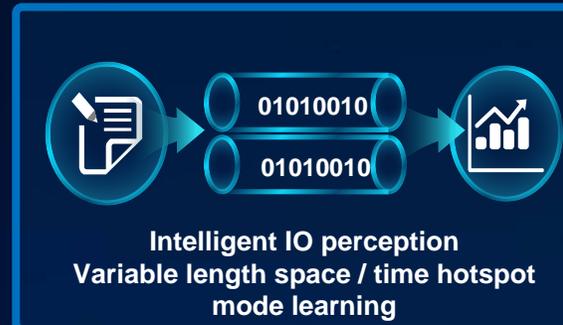
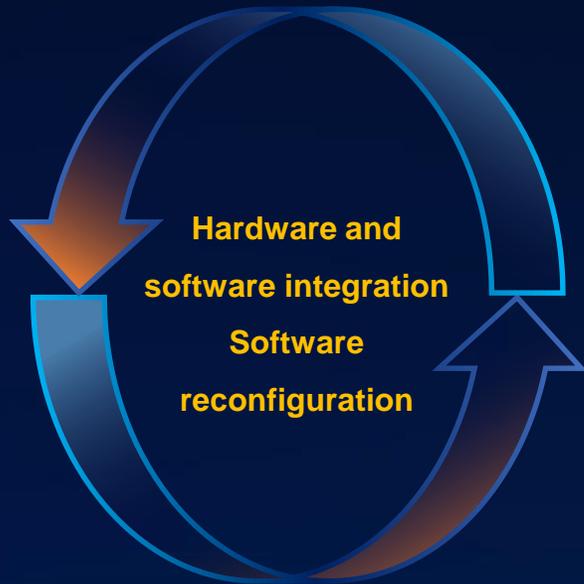


iTurbo2.0 accelerate engine to improve system performance



iTurbo2.0 Intelligent acceleration

All flash new platform



Multivariate computing power

All flash data tiering

Cache delay optimization

Inspur all flash storage performance is in the lead of market



Comparing the performance test results of XX bank customers, Inspur HF18000G5-i is fully ahead of Huawei Dorado 18500 V6

- 8K random read / write (full hit) performance **improved by 21.6%**; Online IO model performance **improved by 16.4%**;
- 8K random write performance **increased by 16.3%**, and 8K random read performance **increased by 17%**

| HF18000G5-I | | | | | | |
|-----------------------------------|------------------------------------|--------------|------------------|--------------------|-----------------|------------------|
| IO model | Model description | IOPS | Bandwidth (MBps) | Average delay (ms) | Read delay (ms) | Write delay (ms) |
| 8K random read / write (full hit) | Limit IO time in full hit scenario | 74,807.91 | 584.44 | 0.132 | 0.101 | 0.153 |
| | | 810,613.34 | 6,332.90 | 0.175 | 0.147 | 0.194 |
| Online IO model | Online IO model | 867,849.42 | 8,135.89 | 0.722 | 0.626 | 0.818 |
| Read batch IO model | Reread batch IO model | 1,635,151.44 | 12,774.62 | 0.774 | 0.761 | 0.897 |
| Double 11 IO model | Double 11 IO model (rewritten) | 468,538.38 | 7,320.91 | 0.771 | 0.323 | 0.963 |
| 8K random write | 8KB random full write | 551,079.65 | 4,305.31 | 0.810 | 0.000 | 0.810 |
| 8K random read | 8KB random read | 1,102,591.10 | 8,613.99 | 0.928 | 0.928 | 0.000 |

| Dorado 18500 V6 | | | | | | |
|-----------------------------------|------------------------------------|------------|------------------|--------------------|-----------------|------------------|
| IO model | Model description | IOPS | Bandwidth (MBps) | Average delay (ms) | Read delay (ms) | Write delay (ms) |
| 8K random read / write (full hit) | Limit IO time in full hit scenario | 61,541.58 | 480.79 | 0.128 | 0.112 | 0.139 |
| | | 488,935.27 | 3,819.81 | 0.162 | 0.121 | 0.189 |
| Online IO model | Online IO model | 745,733.95 | 6,991.44 | 0.852 | 0.890 | 0.814 |
| Read batch IO model | Reread batch IO model | 980,755.52 | 7,662.15 | 0.731 | 0.709 | 0.923 |
| Double 11 IO model | Double 11 IO model (rewritten) | 438,843.69 | 6,856.93 | 0.586 | 0.919 | 0.443 |
| 8K random write | 8KB random full write | 473,661.87 | 3,700.48 | 0.875 | 0.000 | 0.875 |
| 8K random read | 8KB random read | 942,559.15 | 7,363.74 | 0.949 | 0.949 | 0.000 |

Non public test data, please do not transmit

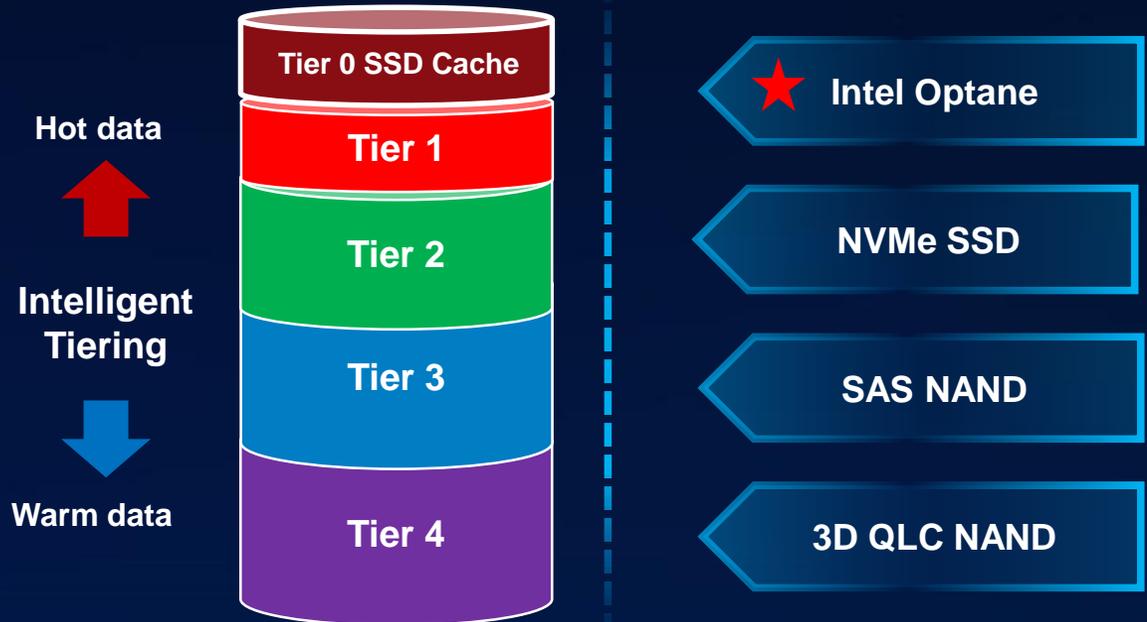
Inspur all flash storage brings a leap in performance through SCM applications



All flash = High performance

Cache acceleration

4+1 layers



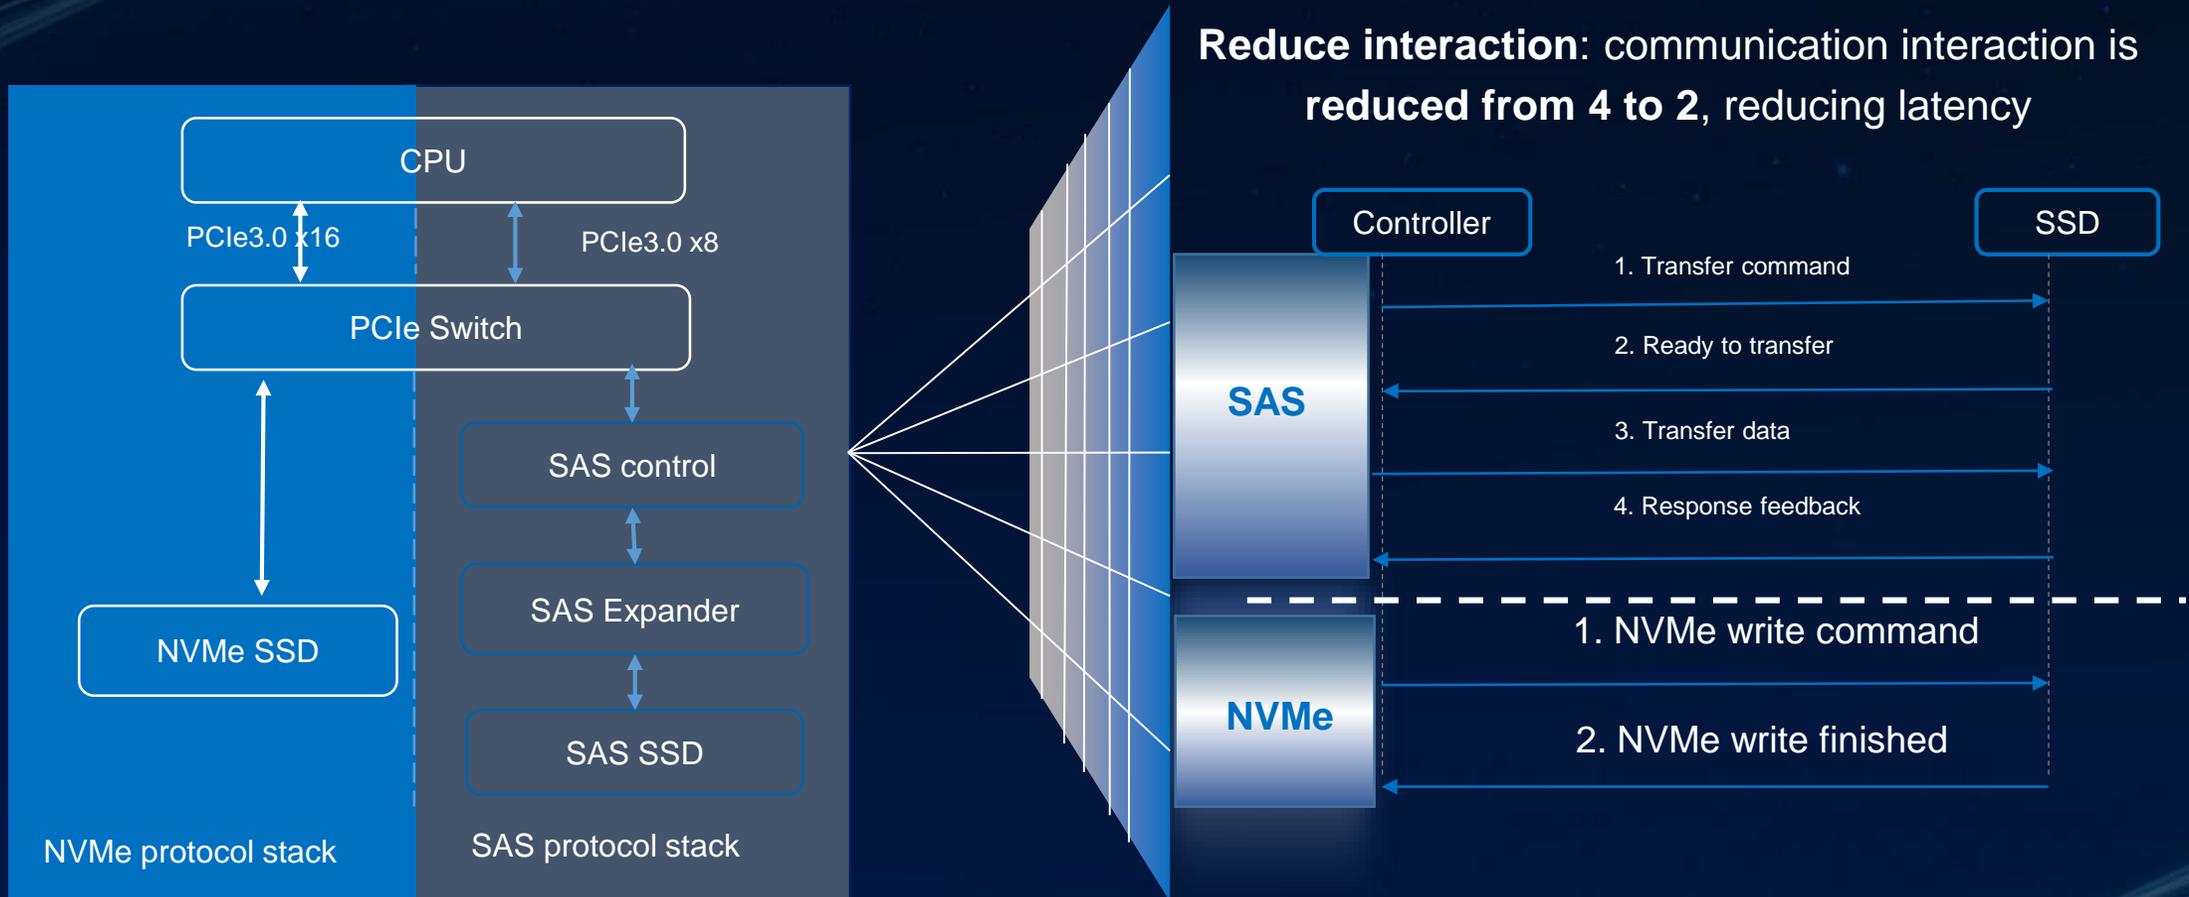
- All flash Tiering, free flow of hot and warm data
- SSD cache, extended cache, accelerate the storage system to effectively deal with application scenarios such as startup storm
- Ensure the business data response performance of core applications
- Data is automatically stored at more cost-effective levels
- Optane improves the average life of the system and reduces the investment in operation and maintenance

Optane all flash=1.35x NVMe all flash=2.5x SAS storage



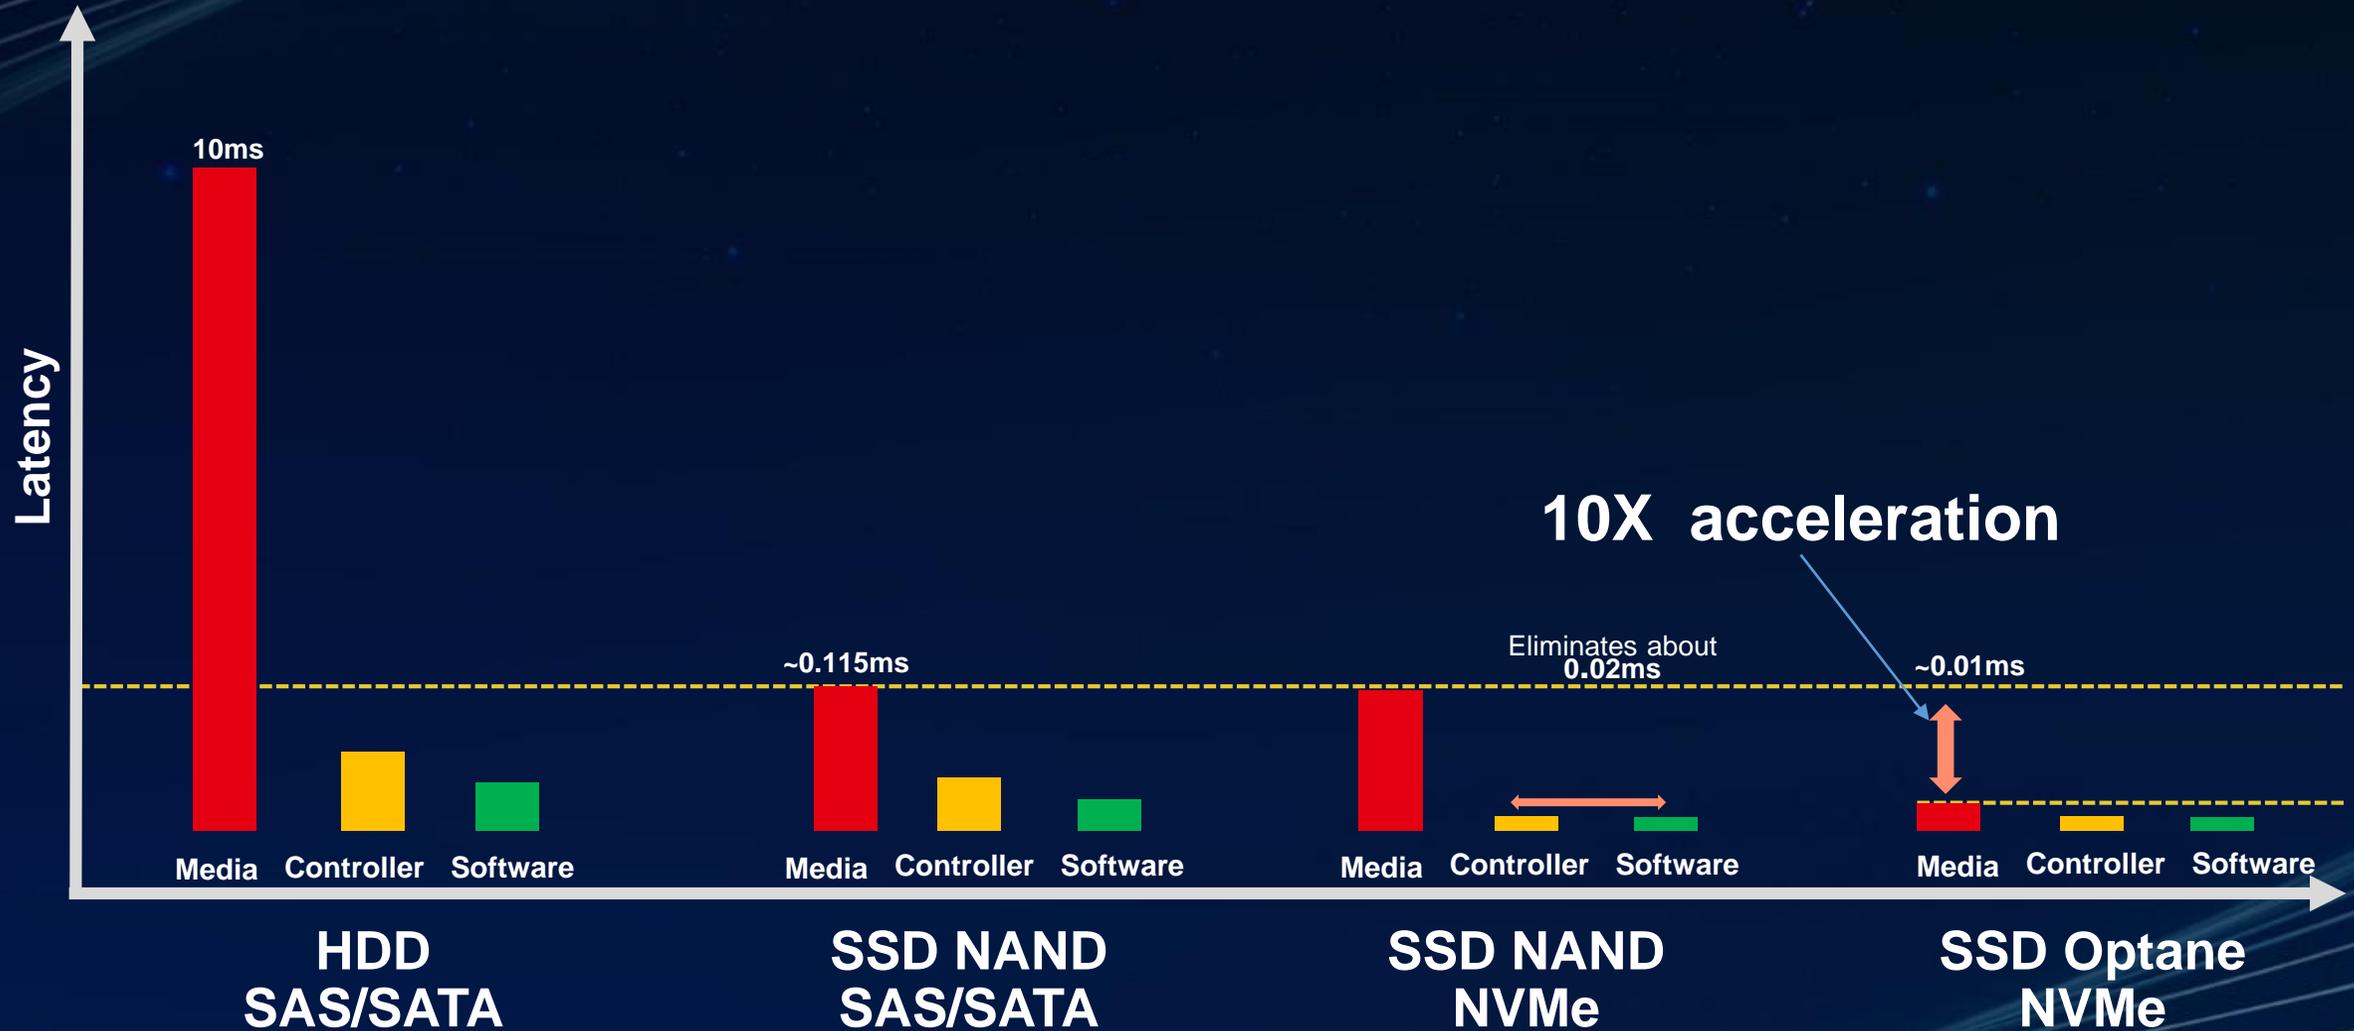
All NVMe protocol reduces processing delay and improves transmission performance

inspur



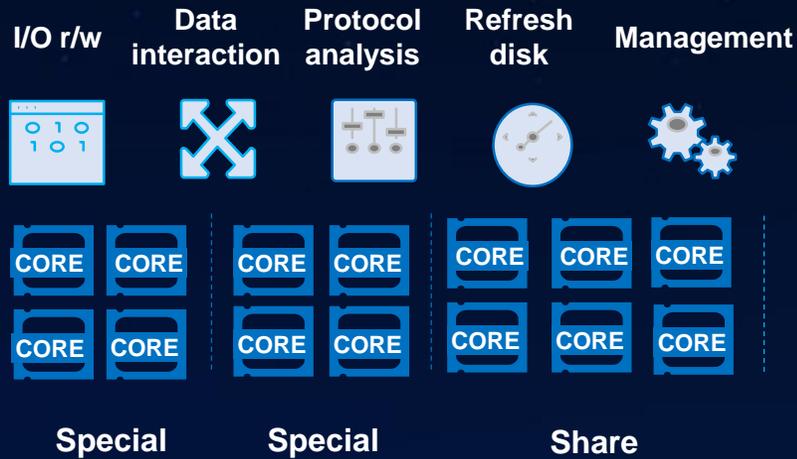
The number of protocol interactions was **reduced from 4 to 2**, the write efficiency was **doubled**, and the number of concurrent transactions was increased to **65536**

Ultimate performance -- the ultimate improvement of optane SSD performance



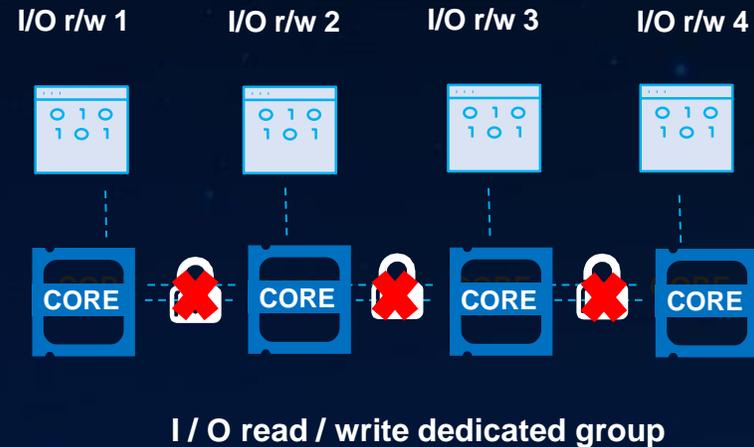
Ultimate performance

-- CPU multi-core load balancing, 20% reduction in latency



CPU split core load balancing

Special core and special purpose to ensure key business performance
Multi core sharing, more balanced CPU operation



Bind IO by core in CPU

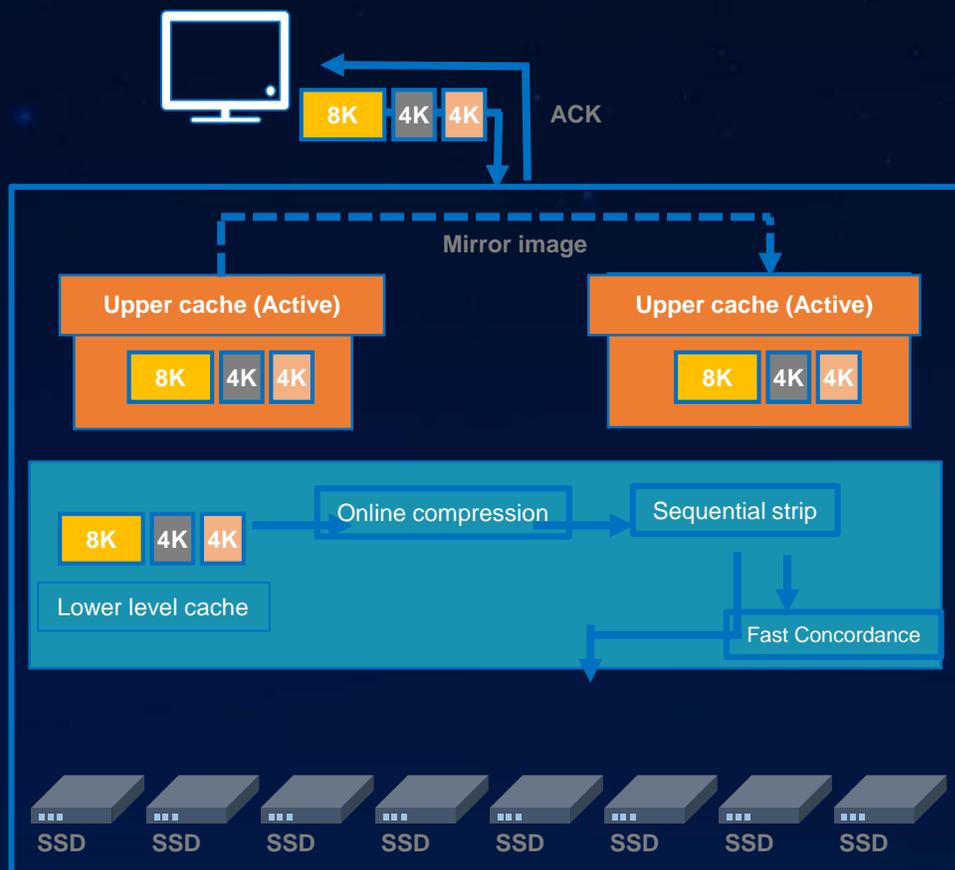
One request is executed continuously on the same core and the lock free design is realized to avoid frequent multi-core switching

CPU core splitting algorithm improves CPU processing efficiency by **2 times**, and the delay is **20% shorter** than that of traditional storage

Ultimate performance

-- cache delay optimization strategy to reduce delay

inspur



1. **Double layer cache design:** reduce the delay of compression / replication by 50%+

2. **Cache partition strategy:** isolate cache resources, ensure key business performance

3. **Cache read ahead algorithm:** read hotspot data in advance, improve cache hit rate and shorten access delay

The cache optimization algorithm reduces the delay by

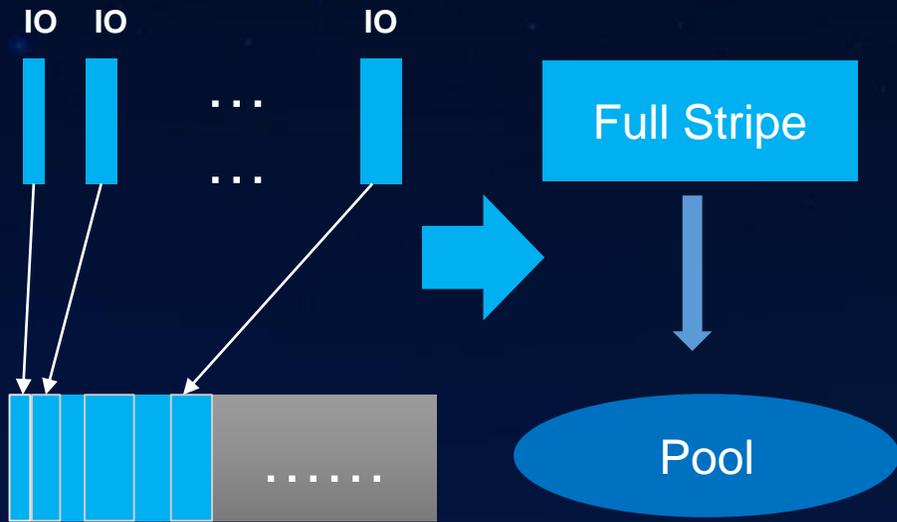
50% and the delay volatility is **< 2%**

Ultimate performance

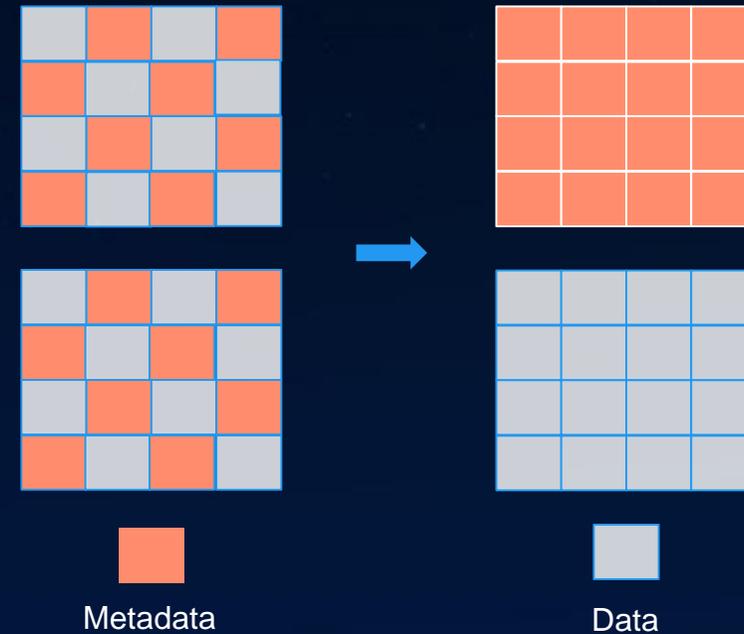
-- full flash optimization algorithm to improve system performance



LSA thin pool



Metadata management



- Multiple discrete data are aggregated into continuous large data blocks, which are written in full stripe order to reduce write amplification and improve performance

- Different types of data synchronously executed in the controller and SSD disk are stored separately

Reduce garbage collection, improve the average service life of SSD and improve system performance

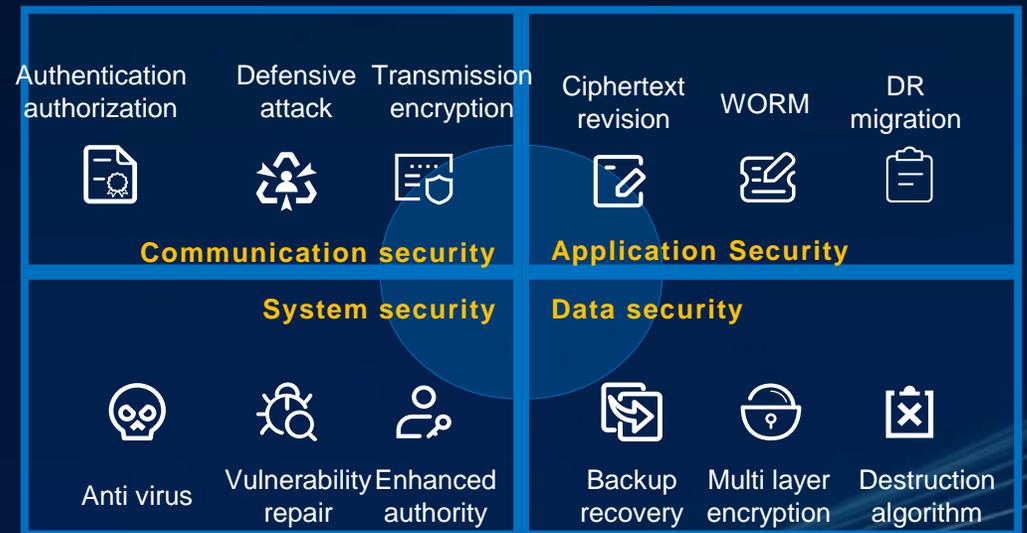
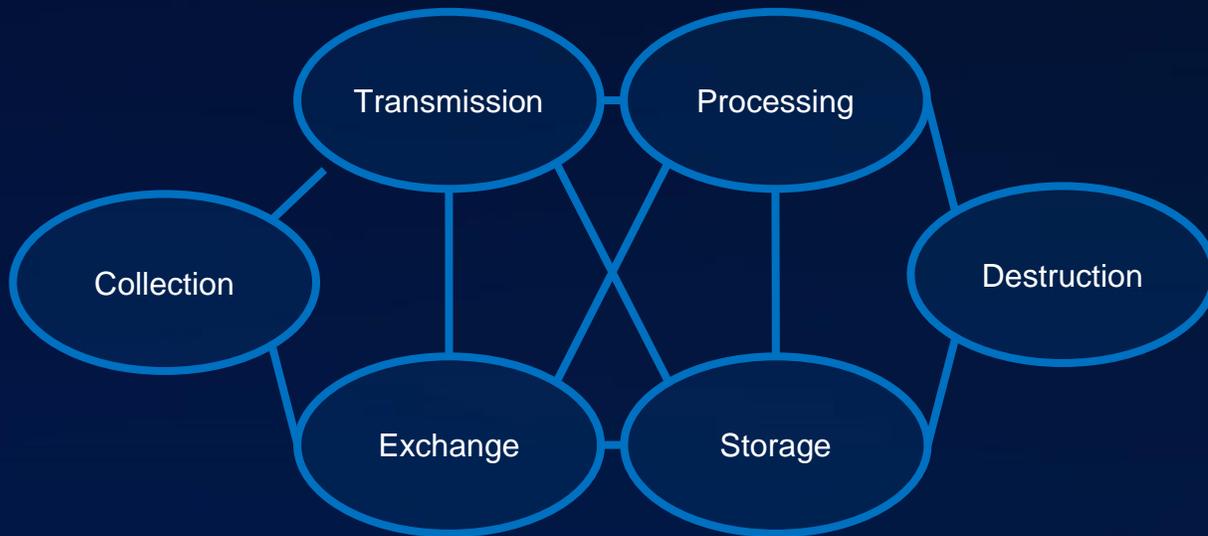
Ultra safe design, achieve full lifecycle data protection



Hold on Take over more than 95% of the data in the industry's storage devices

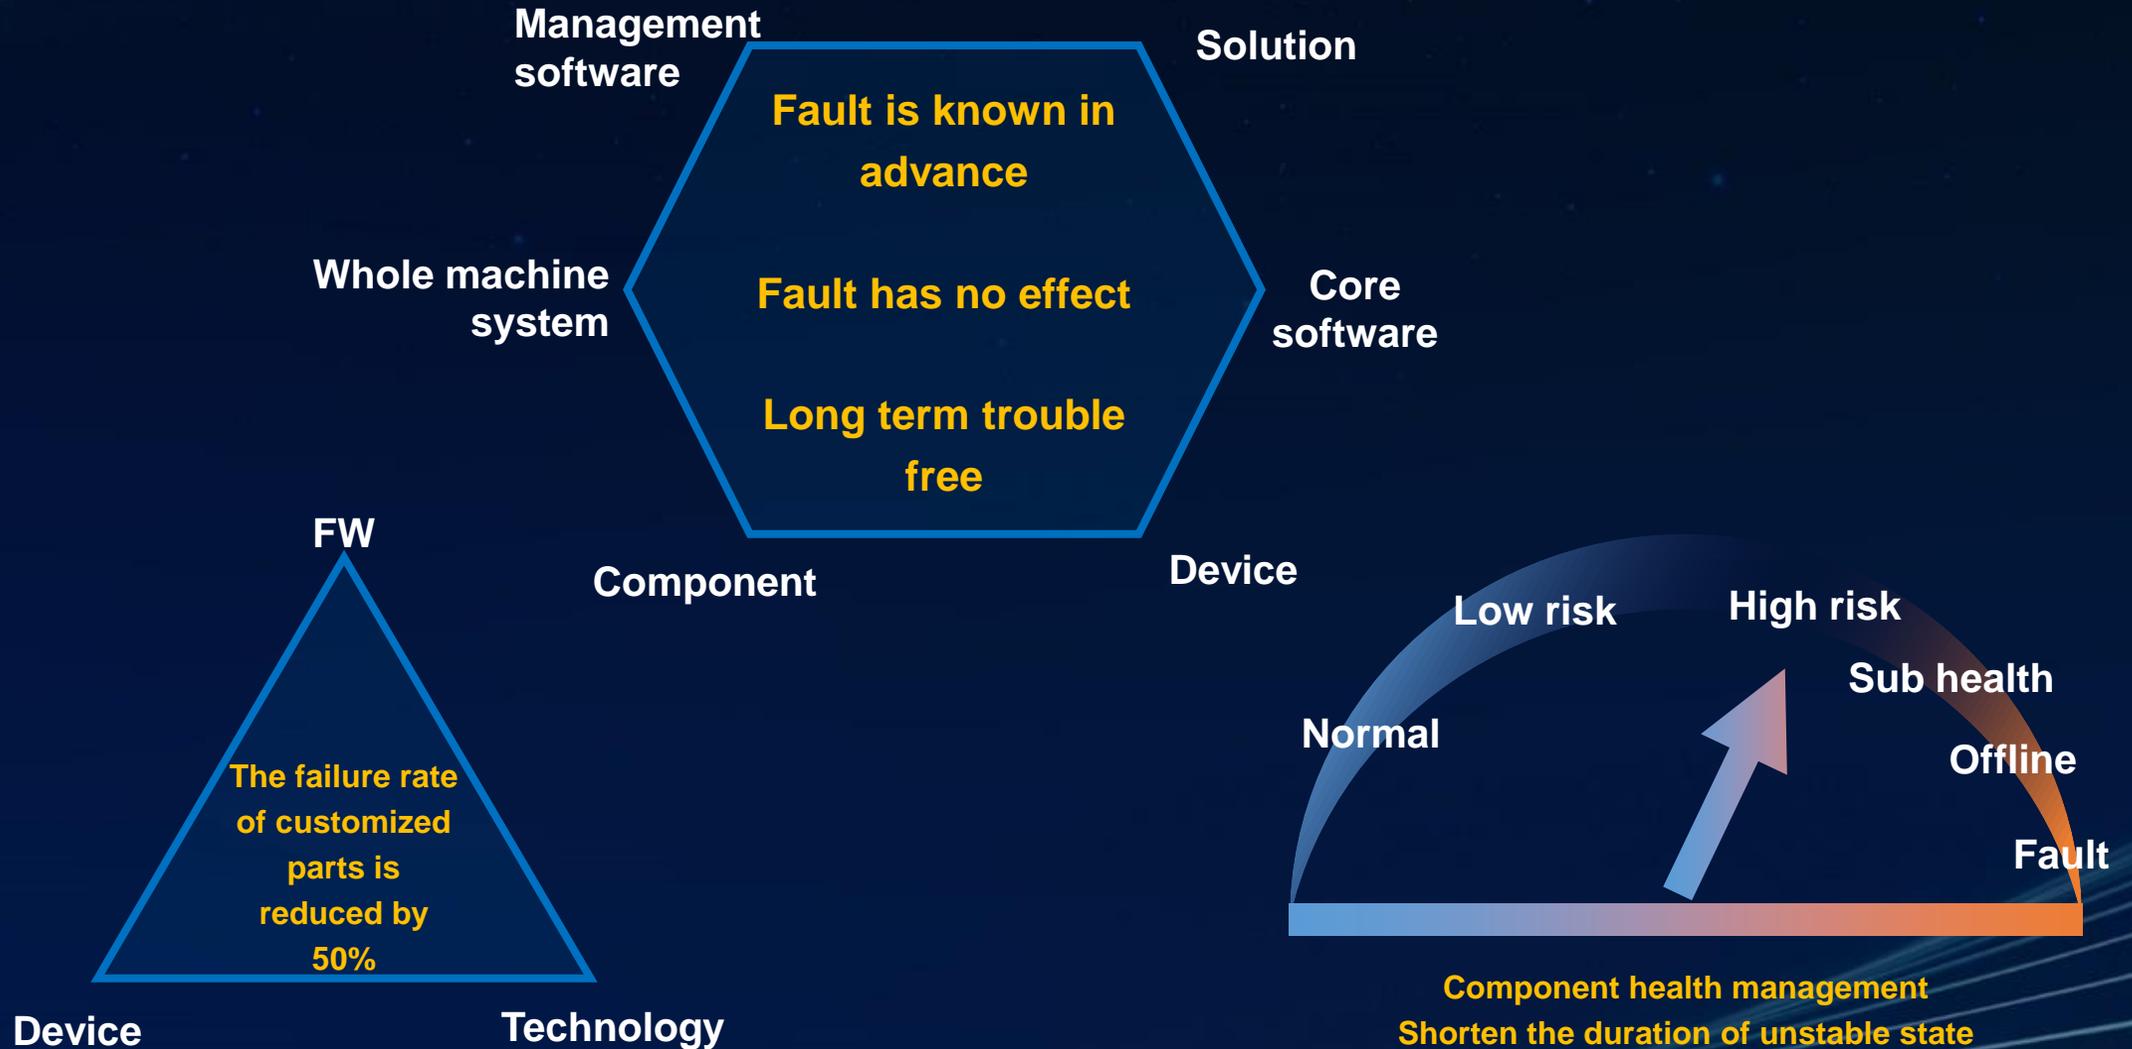
Defensible Clear vulnerabilities according to GBT / gat / ISO standards

Clean up Full media "byte level" destruction



Six layers of protection, business is always online

inspur



Multiple stability design to meet the RAS requirements of core business

inspur



**Component level
design**

Disk fault sector isolation
Full redundancy design of
components
IO module hot plug
Cache secondary standby
power
Global dynamic wear
equilibrium
.....



**System level
design**

Node sub-health
monitoring
Silent data monitoring
Multi control cache copy
Back end media sharing
Raid concurrency
balancing
.....



**Software self
repair data**

Slow disk auto repair
Automatic link repair
BBU automatic monitoring
Data consistency detection
Disk fault prediction
.....



**Reduce the
impact of faults**

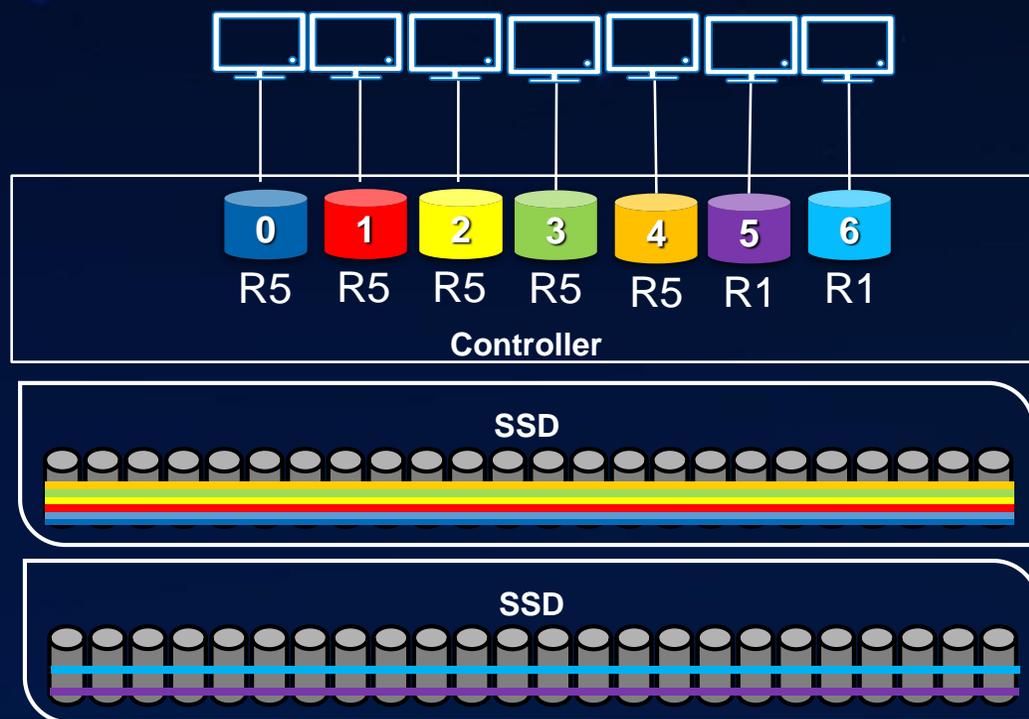
Microcode upgrade does
not affect business
Failover less than 10
seconds
AA service switching within
30 seconds
.....

Extremely reliable InRAID technology

-- global dynamic wear equalization technology

inspur

Greatly improve the service life of SSD
Eliminate your concerns about the service life of SSD



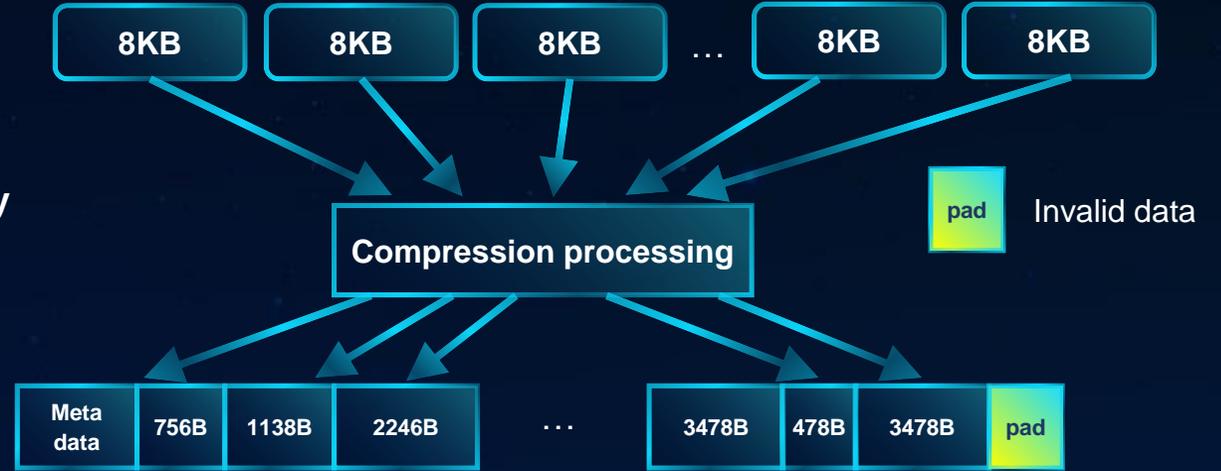
- 1. There is no hot spare concept. Data blocks and hot spare blocks are distributed on all member disks of the RAID array
- 2. The data rebuild speed is greatly reduced, reducing the risk of raid failure in case of disk failure again
- 3. The global stripe mode can avoid the occurrence of SSD hot disks and achieve the average usage of each SSD

Without the concept of hot spare, the hot spare space is evenly distributed, and the data reconstruction speed is up to **10 times**

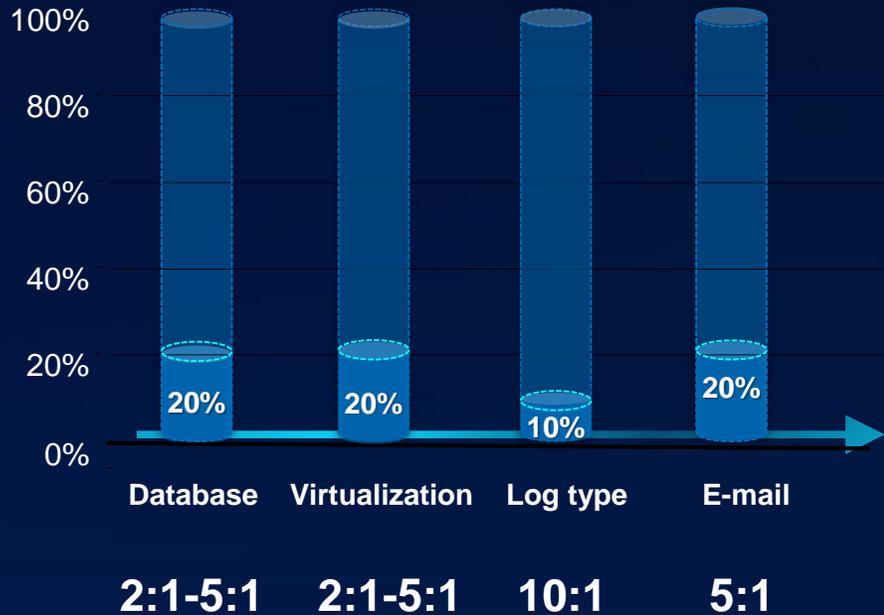
Real time lossless online Duplication & Compression



- Compressed data is 8byte aligned, and the utilization rate of compressed space is 99.8%
- Independent acceleration chip, performance delay lossless
- Weak hash + byte comparison to ensure re deletion efficiency and data consistency



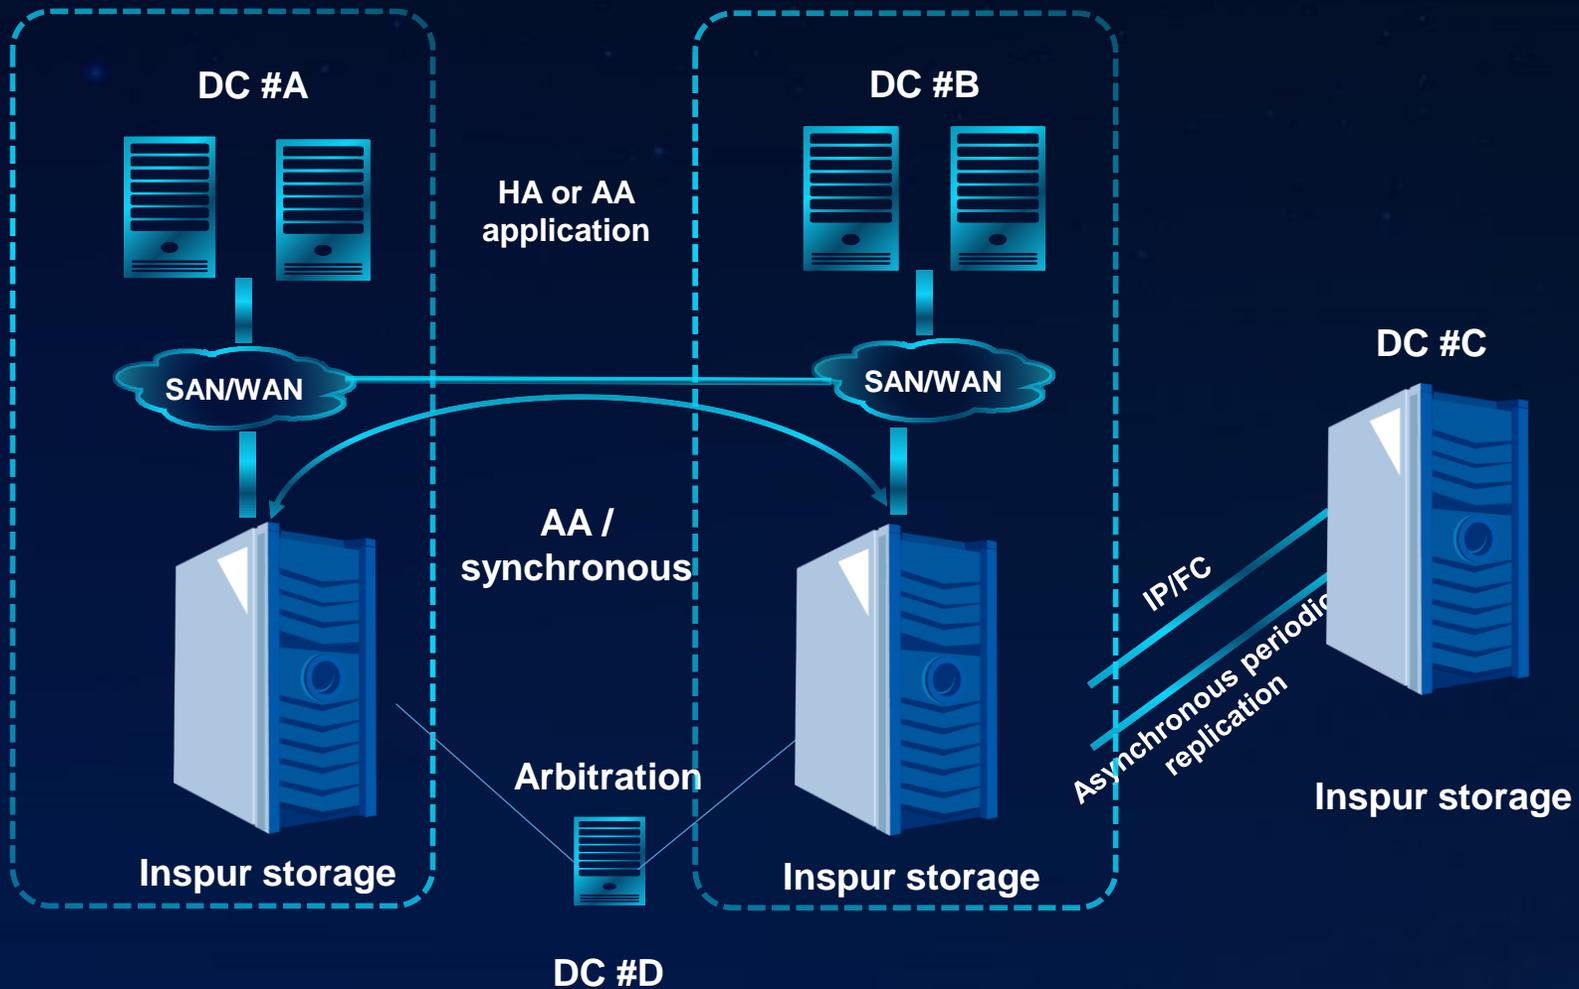
5: 1 Industry leading data reduction efficiency



| Disk space analysis | Alignment granularity (unit: byte) | Spatial complement mathematical expectation (unit: byte) | Proportion of space supplement (complementary length / compressed length) |
|--|------------------------------------|--|---|
| The 8K byte data block is compressed according to the compression experience value of 3:1, The data length range of after compression is 2K ~ 3K bytes | 8 | 4 | 0.133% ~ 0.2% |
| | 1024 | 512 | 16.7% ~ 25% |

Inspur centralized storage 3DC

inspur

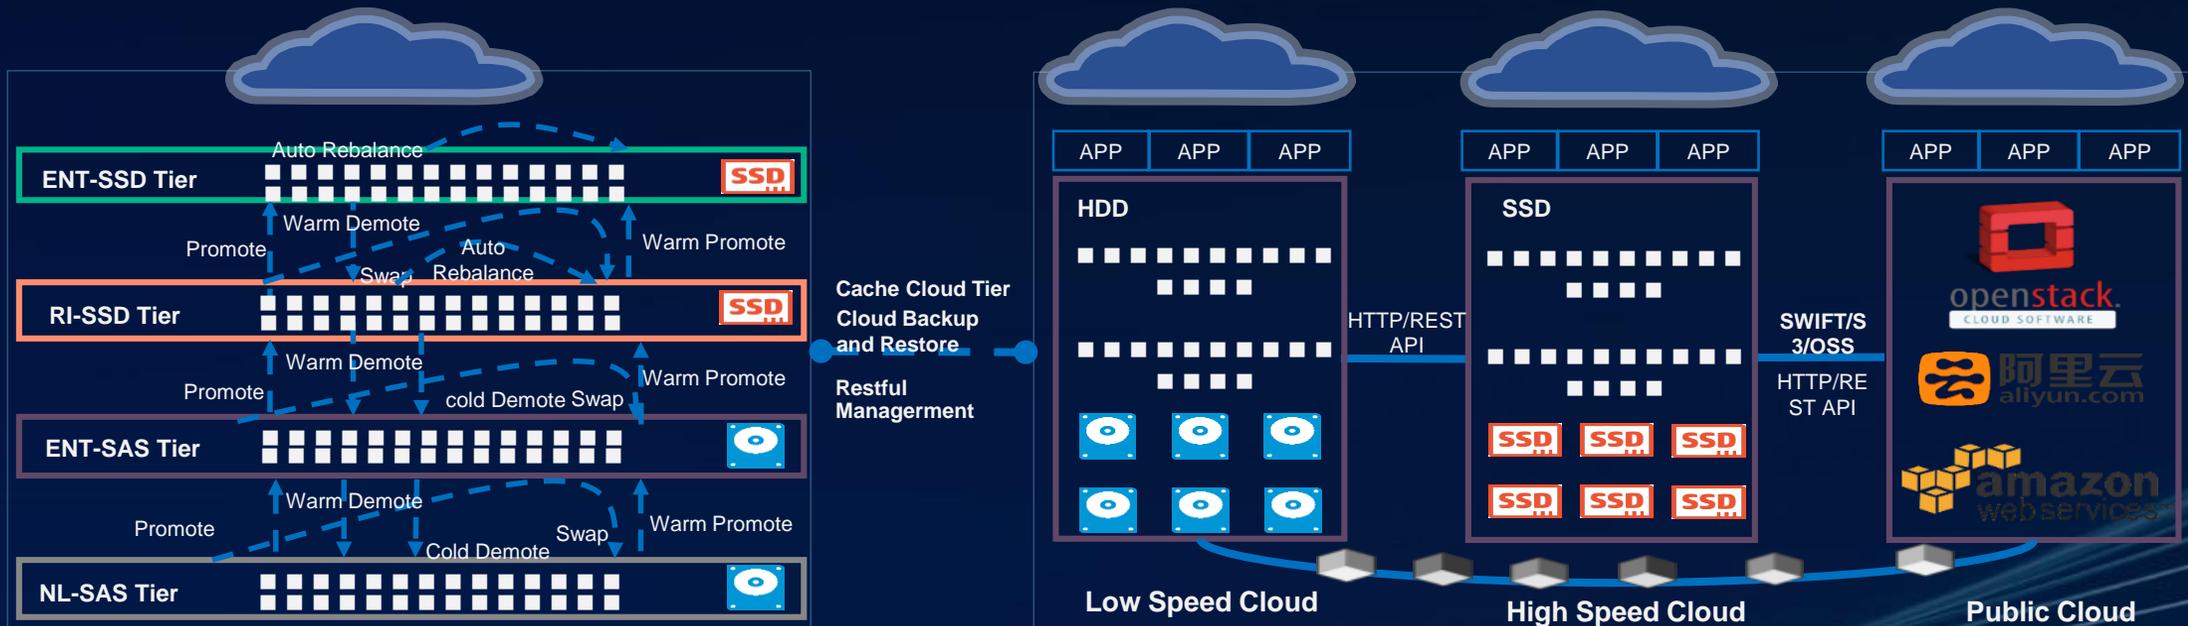


- AA 3DC
- RTO=0/RPO=0
- WAN acceleration X 3

Inspur centralized storage Multiclouds docking



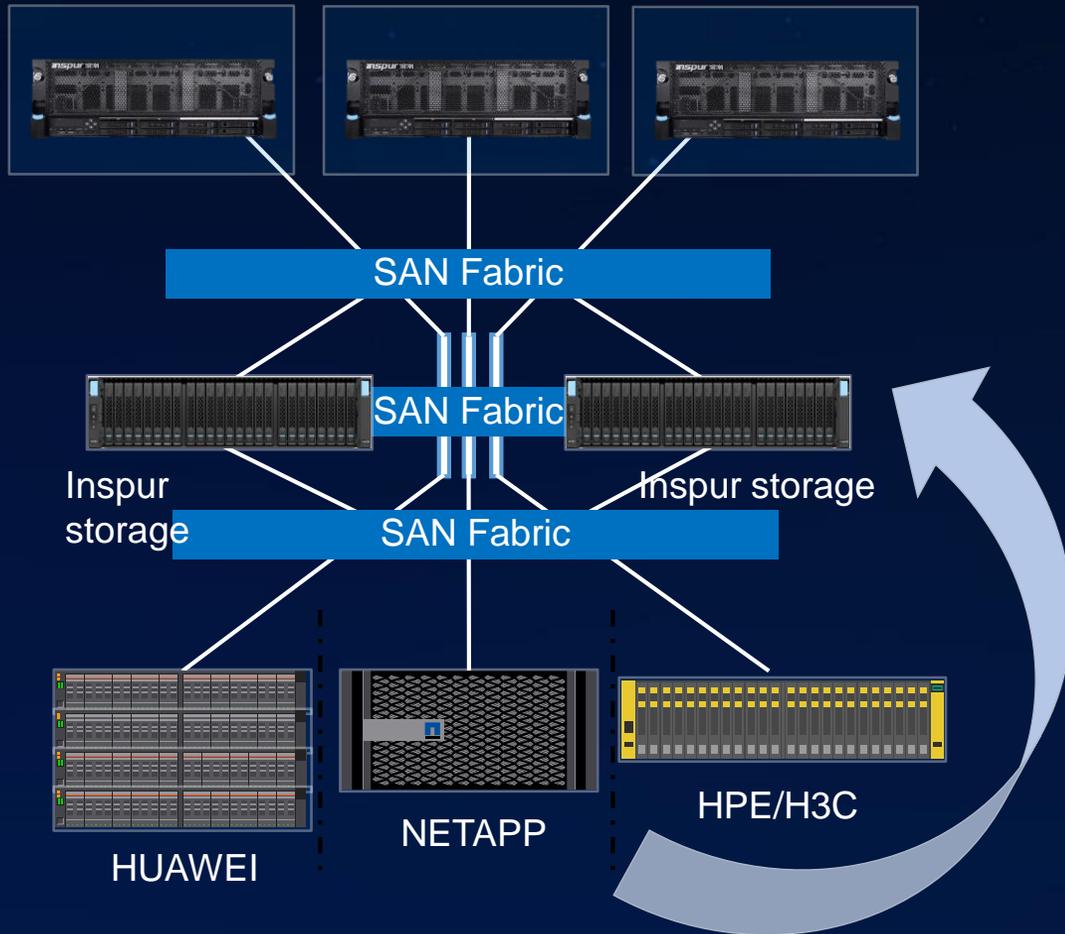
- Cloud data center, Cloud Application
- Full cloud docking, supporting mainstream public and private clouds, cloud backup, cloud archiving, cloud stratification
- Supports data flow between edge data centers, traditional data centers, private clouds, public clouds, multiclouds
- Supports data backup to cloud platforms such as OpenStack, Ali Cloud, AWS, etc.
- Integrated management and scheduling of mixed clouds and multiclouds, putting the right data in the right place at the right time, to achieve the best performance-price ratio



Inspur Centralized Storage

Allows heterogeneous integration of more than 95% of industry storage models

inspur



Application Requirements

- Old Storage can not satisfy application and needs to be replaced and reuse
- Resources between old and new storage systems cannot be integrated and managed
- Business continuity needs to be guaranteed during data migration

Solution value

- Compatible with heterogeneous storage and simplified storage management
- Integrate equipment space and simplify space allocation
- Reuse old storage and make full use of existing resources

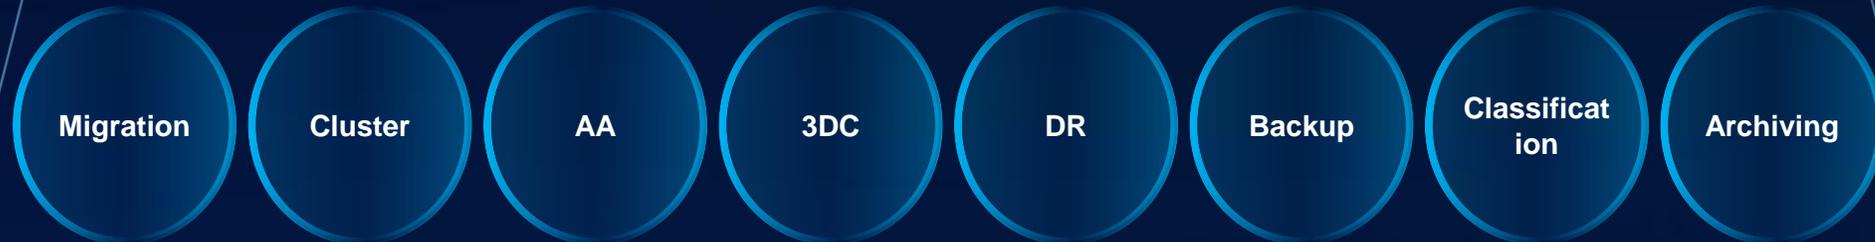
New way of storage -- Data lifecycle storage infrastructure

inspur

Business security, efficiency and continuity, data lifecycle management



Data "system" storage infrastructure



Edge storage

Centralized storage
platform

Distributed storage
platform

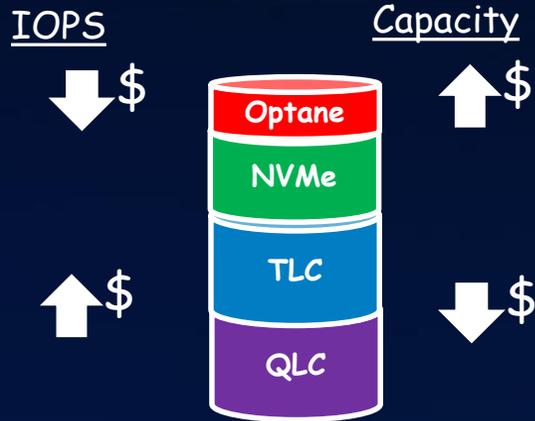
Storage Gateway

Inspur all flash storage helps build a new cloud data center

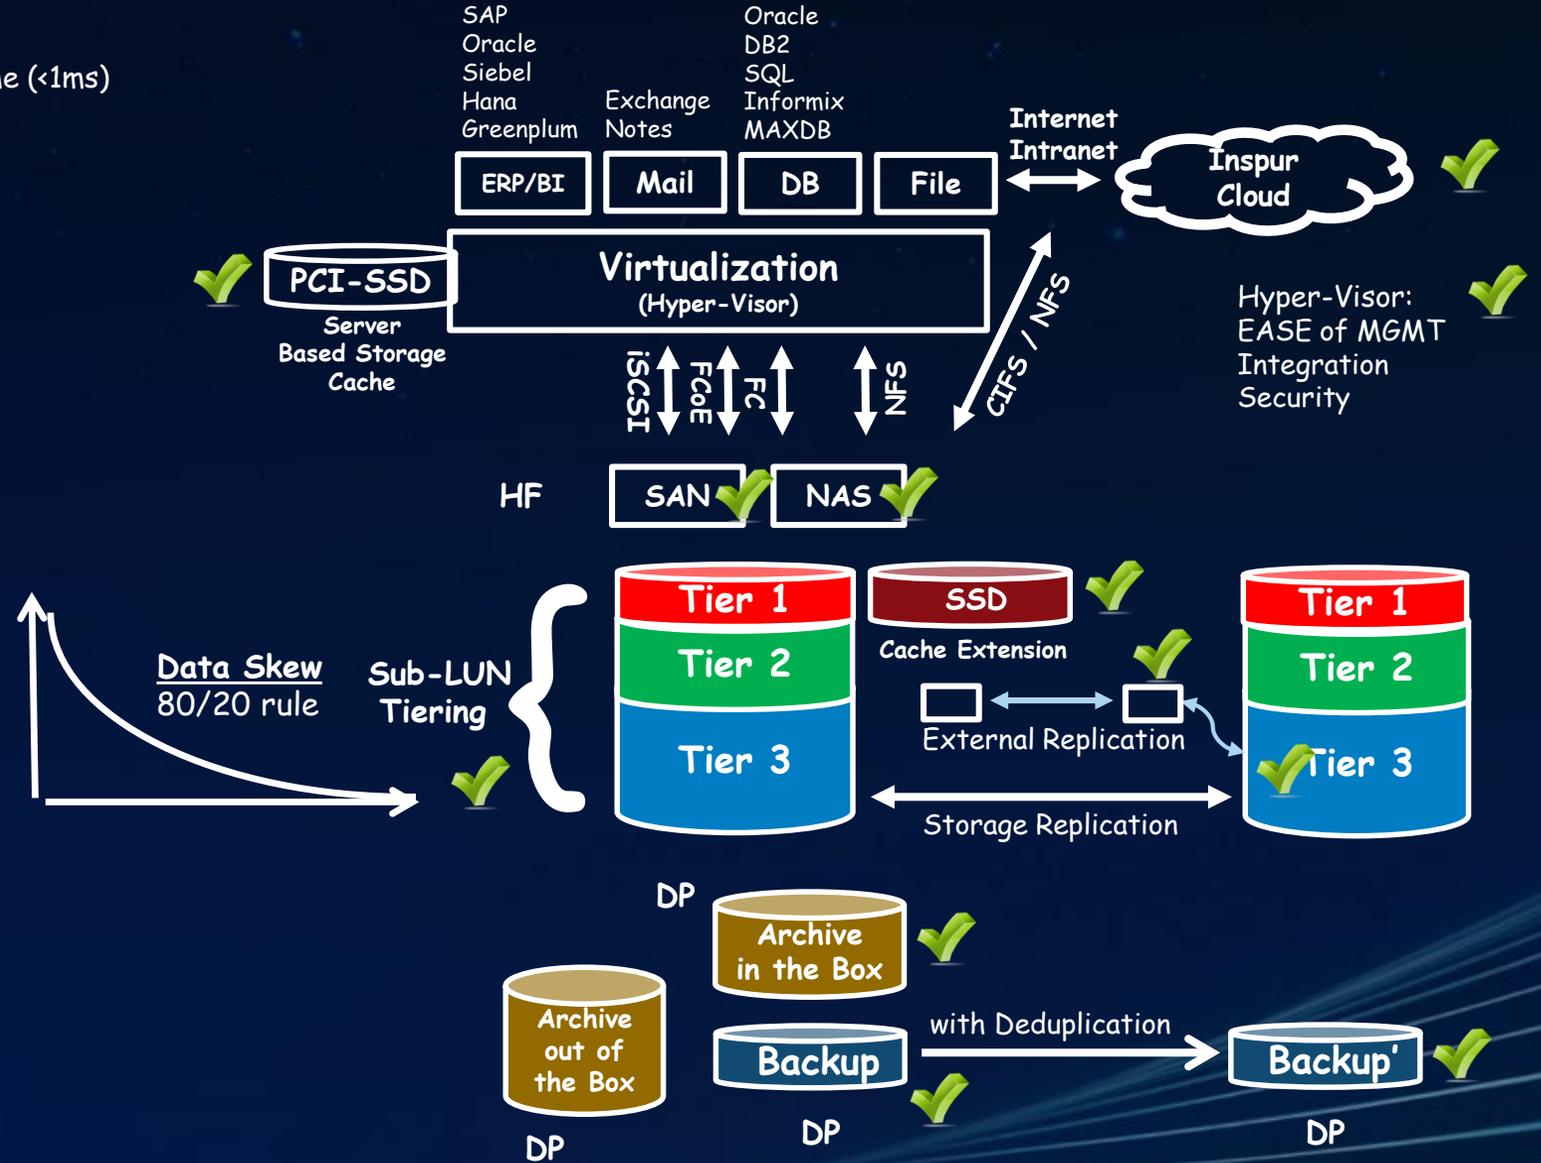


Value of SSD
 Reduce application response time (<1ms)
 Reduce environment costs

- Power
- Cooling
- Floor space

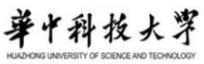
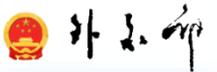


- Storage Function:**
- Snap / Clone
 - Continuous Data Protection
 - Synchron / Asynchron Replication
 - File Deduplication
 - Block Compression
 - Thin / Thick Devices
 - Spin Down
 - Worm Function
 - Centralized Management



Inspur all flash storage has gradually become the preferred solution for customers in the industry

inspur

| Finance | Government | Medical treatment | Energy sources | Communication | Traffic | Education | Manufacture |
|---|---|---|---|---|---|---|---|
|  中国工商银行 INDUSTRIAL AND COMMERCIAL BANK OF CHINA |  中华人民共和国工业和信息化部 Ministry of Industry and Information Technology of the People's Republic of China |  |  SINOPEC |  中国移动 China Mobile |  中国国家铁路集团有限公司 CHINA RAILWAY |  华中科技大学 HUAZHONG UNIVERSITY OF SCIENCE AND TECHNOLOGY |  广汽集团 GAC GROUP |
|  中国銀行 BANK OF CHINA |  中华人民共和国农业农村部 Ministry of Agriculture and Rural Affairs of the People's Republic of China |  山东大学齐鲁医院 QILU HOSPITAL OF SHANDONG UNIVERSITY |  国家电网有限公司 STATE GRID CORPORATION OF CHINA |  中国电信 CHINA TELECOM 世界触手可及 |  中国民航 CAAC |  北京航空航天大学 BEIHANG UNIVERSITY |  ANSTEEL 鞍钢集团 |
|  中国邮政储蓄银行 POSTAL SAVINGS BANK OF CHINA |  外交部 |  郑州大学第一附属医院 The First Affiliated Hospital of Zhengzhou University 郑州大学第一临床学院 (原河南医科大学第一附属医院) |  中国石油 |  China unicom 中国联通 |  中国東方航空 CHINA EASTERN |  山东第一医科大学 Shandong First Medical University & Shandong First University of Medical Sciences |  LUXSHARE ICT 立讯精密 |
|  中国人民银行 THE PEOPLE'S BANK OF CHINA |  国家税务总局 State Taxation Administration |  武汉大学中南医院 ZHONGNAN HOSPITAL OF WUHAN UNIVERSITY |  国家能源集团 CHN ENERGY 中国神华能源股份有限公司 CHINA SHENHUA ENERGY COMPANY LIMITED |  中国通信服务 CHINA COMSERVICE |  河北高速公路集团有限公司 HEBEI EXPRESSWAY GROUP LIMITED |  北京科技大学 University of Science and Technology Beijing |  中国烟草 CHINA TOBACCO |
|  交通銀行 BANK OF COMMUNICATIONS |  中国海关 CHINA CUSTOMS |  中国科学技术大学附属第一医院 THE FIRST AFFILIATED HOSPITAL OF USTC 安徽省立医院 ANHUI PROVINCIAL HOSPITAL |  中国南方电网 CHINA SOUTHERN POWER GRID |  天翼云 Chinatelecom Cloud |  广州港集团 GUANGZHOU PORT GROUP |  中国科学院 CHINESE ACADEMY OF SCIENCES |  长城汽车 Great Wall Motors |

Inspur all flash storage product specifications



| Product name | HF5000G5-MS25 (SAS) | HF5000G5-HS25 (SAS) | HF5000G5-MN25 (NVMe) | HF6000G5-M (SAS & NVMe) | HF6000G5-H (SAS & NVMe) | HF8000G5 (SAS & NVMe) | HF18000G5 (SAS & NVMe) |
|---|---|------------------------|-------------------------|---|----------------------------|--------------------------|---------------------------|
| Controller cabinet form | 2U25 | 2U25 | 2U25 | 4U independent handpiece | 4U independent handpiece | 4U independent handpiece | 4U independent handpiece |
| Controller QTY | 2~16 | 2~16 | 2~16 | 2~16 | 2~16 | 2~32 | 2~32 |
| Processor (per dual controller) | 2*10 core | 2*16 core | 2*16 core | 4*16 core | 4*20 core | 4*20 core | 4*20 core |
| Cache capacity (per dual controller) | 128GB/256GB | 256GB/512GB/1TB | 256GB/512GB/1TB | 512GB/1TB/2TB | 1TB/2TB/3TB | 1TB/2TB/3TB | 1TB/2TB/3TB |
| Host interface card (per dual controller) | 6 | 6 | 6 | 24 | 24 | 24 | 24 |
| Disk channel interface | SAS3.0 | SAS3.0 | SAS3.0/PCIe3.0 | SAS3.0/PCIe3.0 | SAS3.0/PCIe3.0 | SAS3.0/PCIe3.0 | SAS3.0/PCIe3.0 |
| Optane | N | | Y | Y | Y | Y | Y |
| SSD type | SAS SSD | | SAS/NVMe SSD | SAS/NVMe SSD | SAS/NVMe SSD | SAS/NVMe SSD | SAS/NVMe SSD |
| Maximum number of hard disks | 1200 | 1400 | 1500 | 2000 | 2400 | 3400 | 4800 |
| Compression | Onboard compression chip (2 engines / controller) | | | Onboard compression chip (3 engines / controller) | | | |
| ACC card | N | | | YY | | | |
| IO card type | 1/10/40Gb ISCSI, 16/32Gb FC | | | | | | |

inspur

Thank you!